

Geospatial-Temporal Semantic Graphs for Remote Sensing Data Analysis



Mark Rintoul
Aug., 2015



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Team

Tri-Lab PIs:

- Randy Roberts, LLNL
- Randy Brost, SNL
- Paul Pope, LANL
- Danny Rintoul, SNL

Universities:

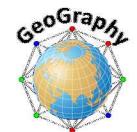
- Jarlath O-Neil-Dunne, UVM
- Sean McFaden, UVM
- Taylor Engel, UVM
- Dave Messinger, RIT

Support:

- DOE NA-22 Office of Non-Proliferation.
- Sandia LDRD Office.

Team:

- Craig Blackhart, LANL
- Michelle Carroll, SNL
- Debbie Dennison, LLNL
- Luci Gaines, LANL
- John Goforth, LLNL
- Art Jolly, LLNL
- Will McLendon, SNL
- Dan Morrow, SNL
- Jo Ellen Neuman, LLNL
- Ojas Parekh, SNL
- Drew Patterson, SNL
- David Perkins, SNL
- Lakshman Prasad, LANL
- Jorge Roman, LANL
- Shelly Spearing, LANL
- Diane Woodbridge, SNL



Overview



- Motivation.
- Computation.
- Example Results.
- Extensions.
- Discussion.

Motivation



Motivation:

- Overwhelming remote sensing data.

General approach:

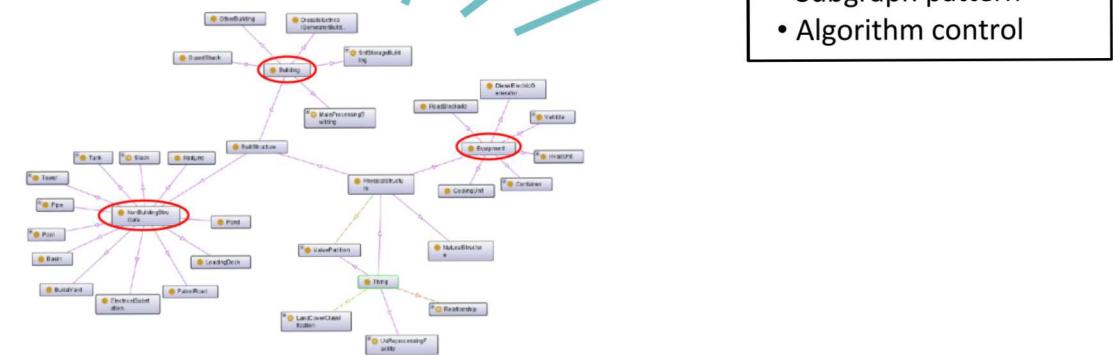
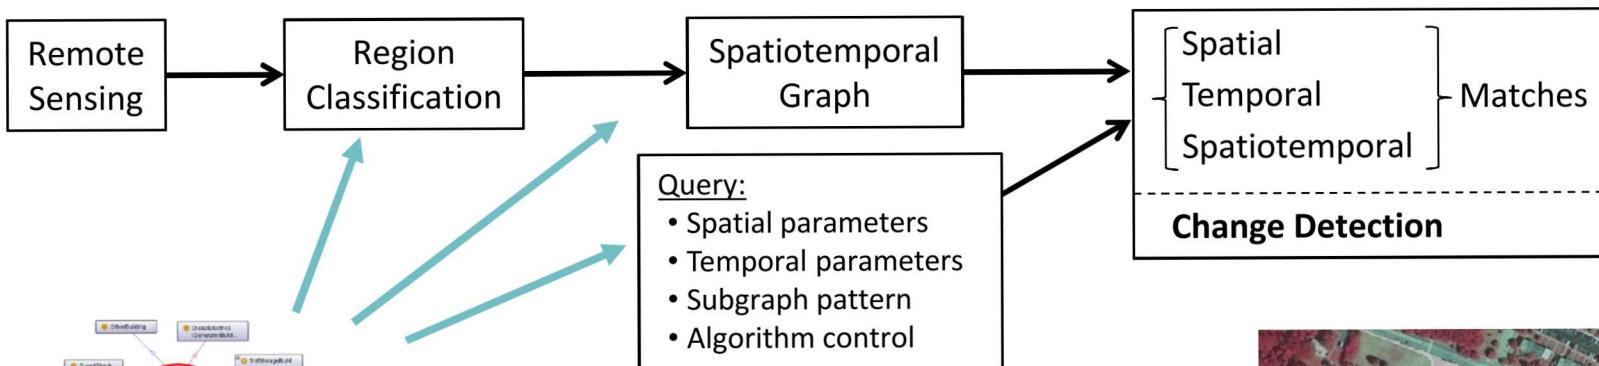
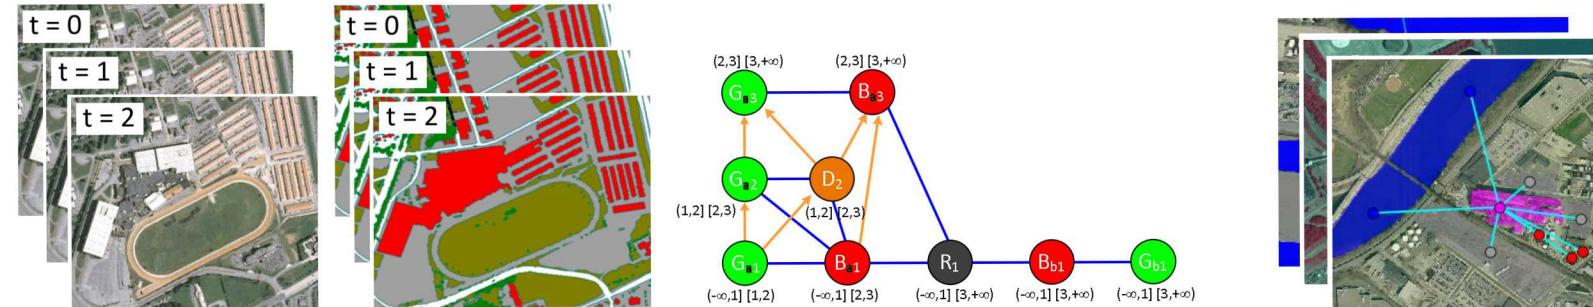
- Automatically find items of potential interest.
- “Cue” to user for review.

Types of questions:

- Geospatial: Find all power plants.
- Temporal: Find all changes.
- Geospatial-temporal: Find all power plants that changed.
- Multi-modality: Find new construction near points of interest.
Find industrial facilities with unusual emissions.

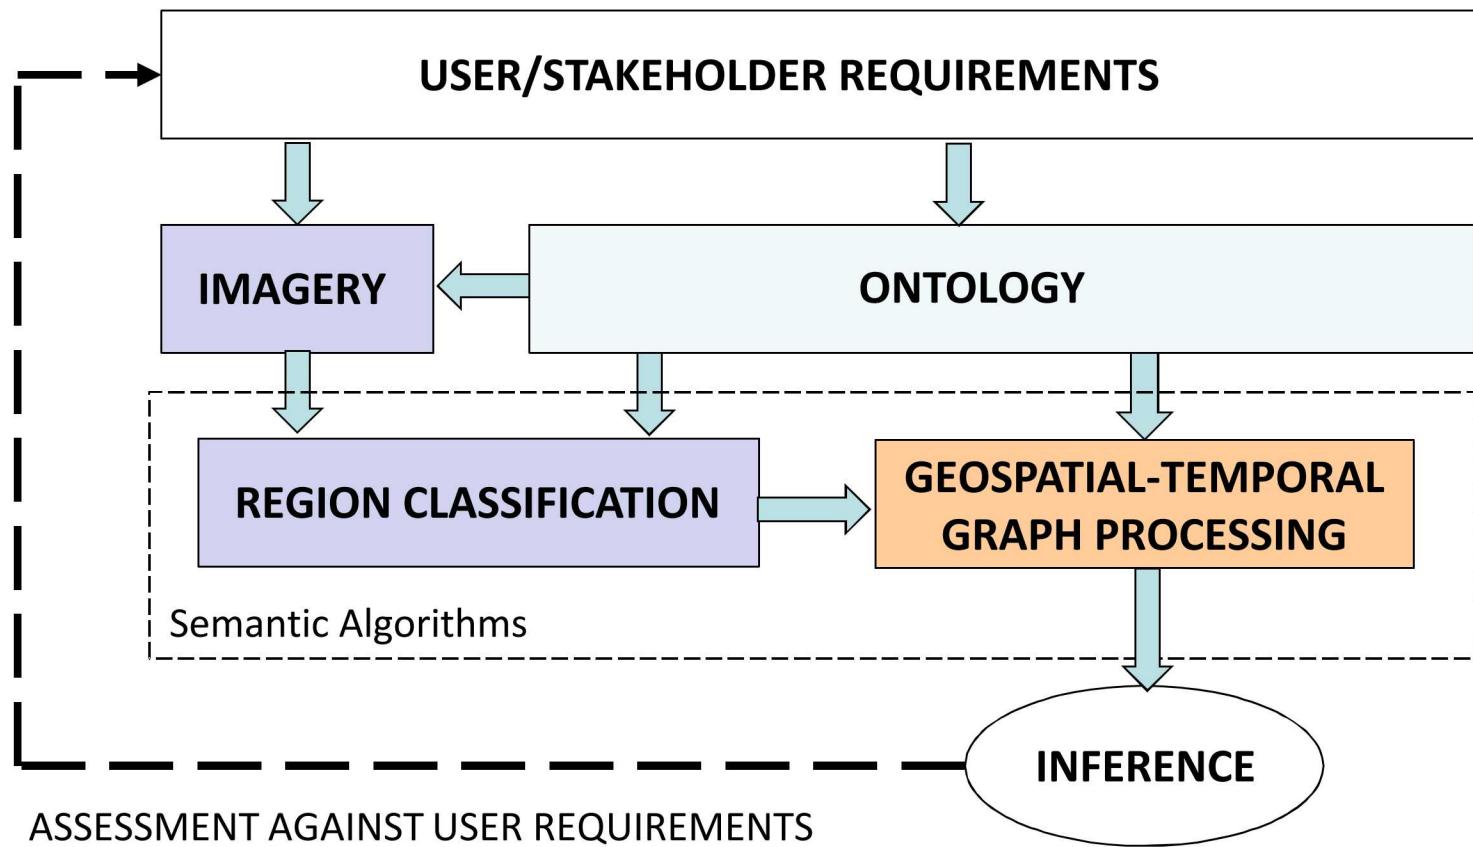
In sum: Make remote sensing data searchable, over space, modality, and time.

Geospatial-Temporal Graph Data Flow



Ontology

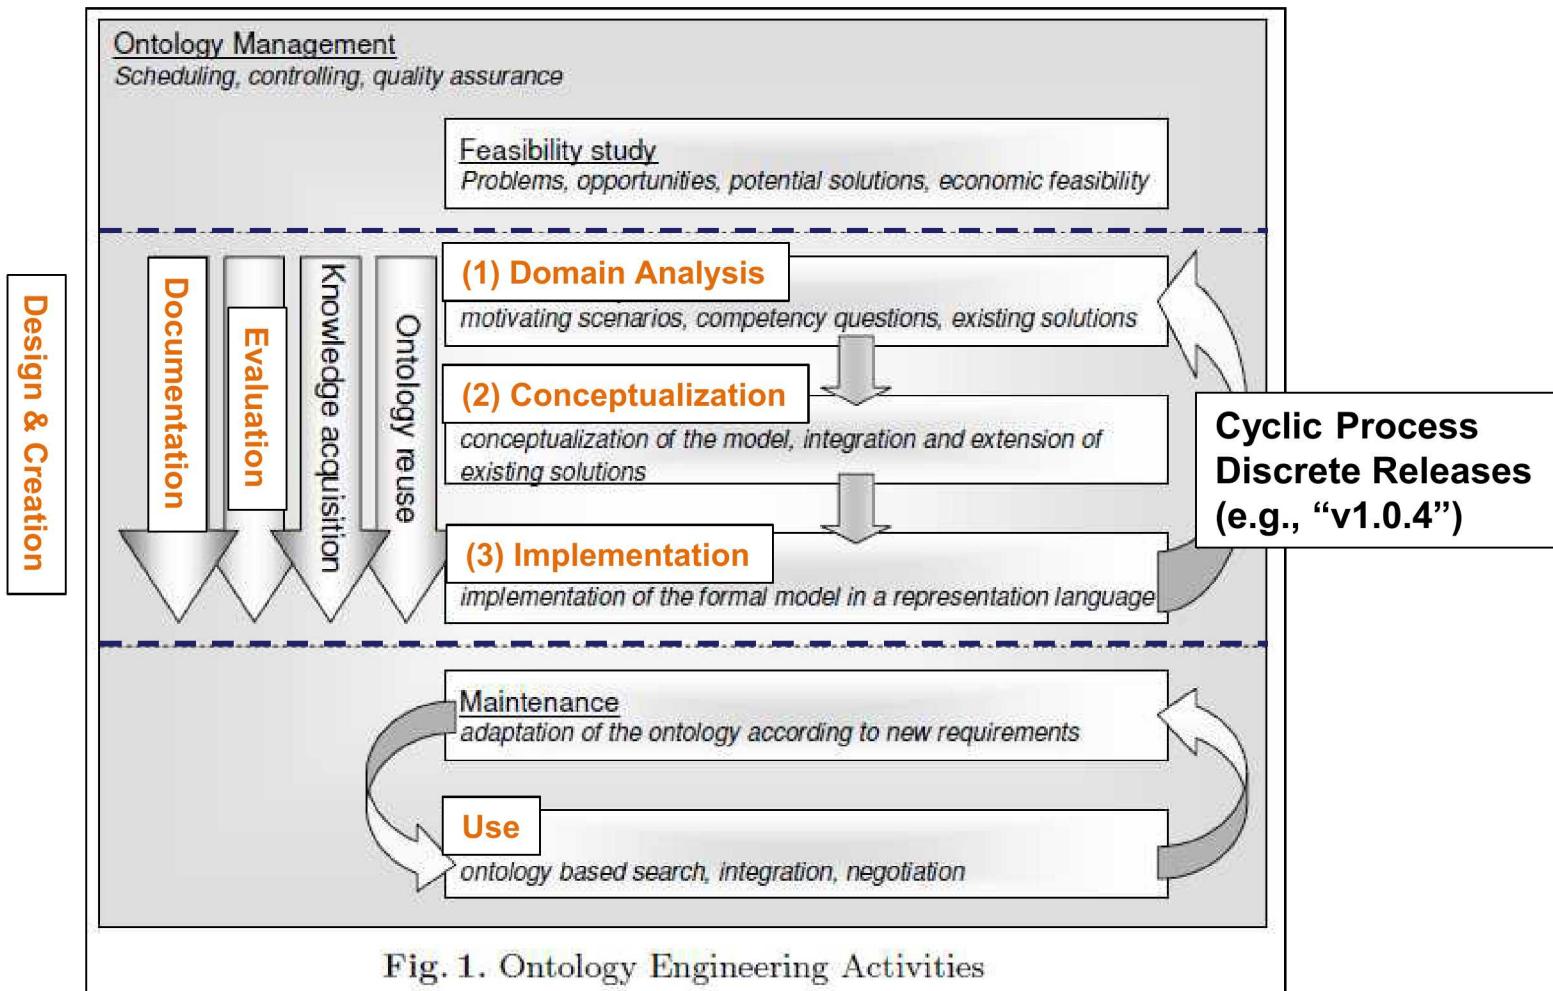
Tri-Lab Project Overview



- ALL
- Los Alamos (LANL)
- Lawrence Livermore (LLNL)
- Sandia (SNL)

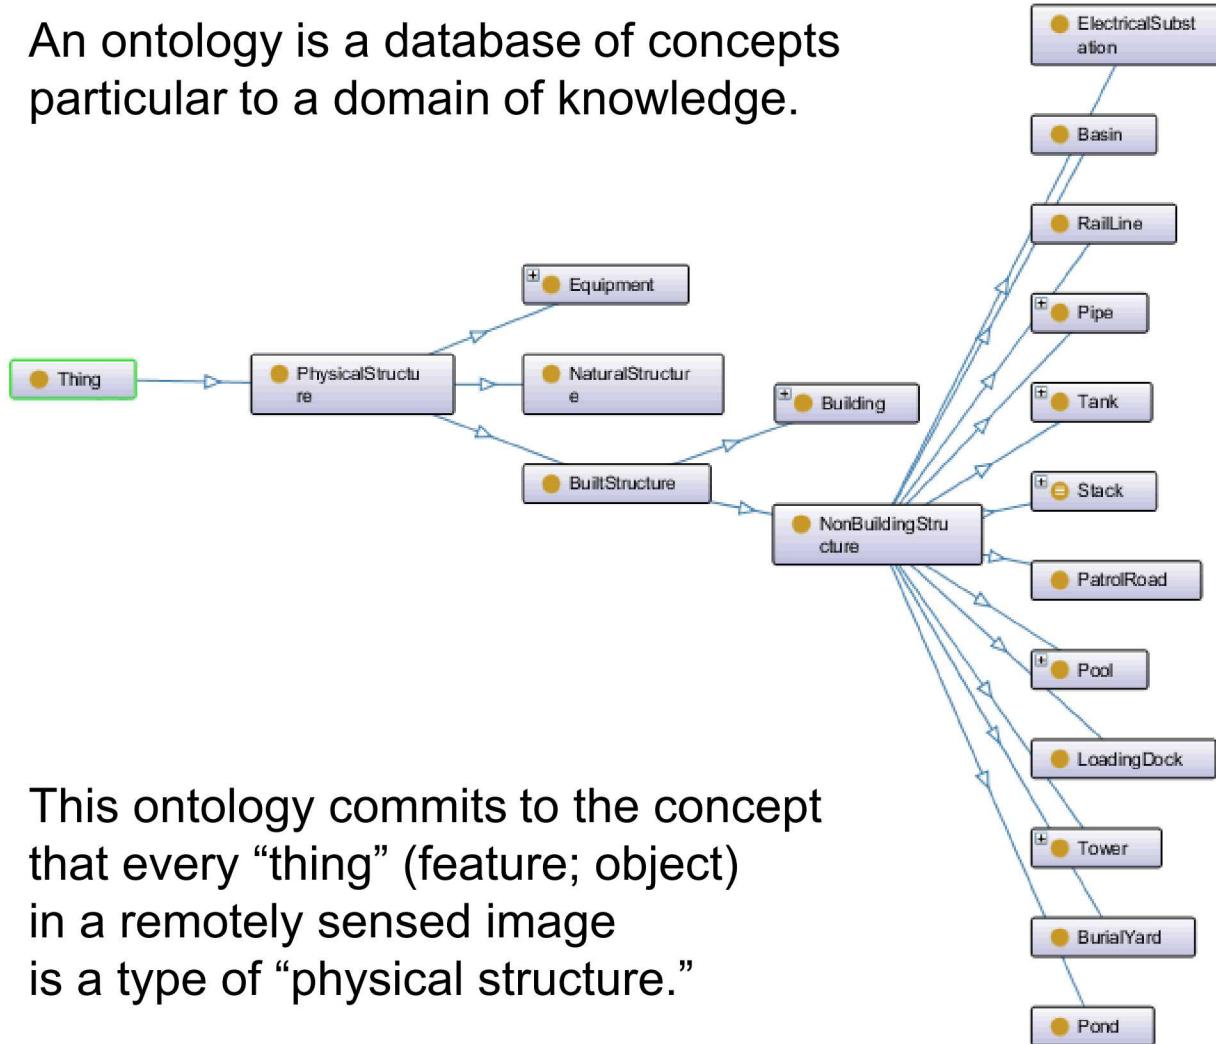
Funded by the DOE NA-22 Office of Non-Proliferation.

Ontology Engineering



Ontology

An ontology is a database of concepts particular to a domain of knowledge.



This ontology commits to the concept that every “thing” (feature; object) in a remotely sensed image is a type of “physical structure.”

Benchmark Imagery Suite

A task-specific test data set:

- VIS and VIS/NIR aerial images:
 - Chemical Processing (27) — refineries, ...
 - Heavy Manufacturing (31) — steel foundries, ...
 - Heat Processing (37) — power plants, ...
 - Mechanical Processing (83) — aluminum processing, ...
 - Semiconductor (12) — chip companies, ...
- Synthetic images (38 fictitious electrical power plants).
- LiDAR elevation rasters (16 of the real facilities).

Benchmark Image Examples

Image properties:

- Orthorectified GeoTIFF.
- Near-nadir viewing angle.
- Area: 1 to 1.6 km².
- GSD: 0.3 to 1m.

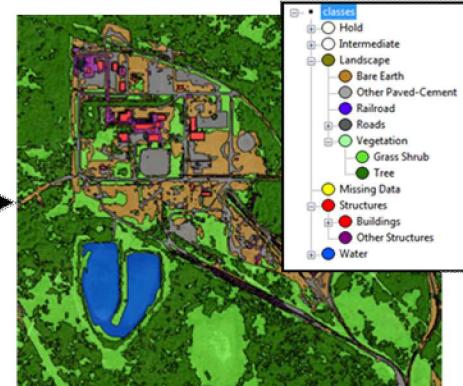
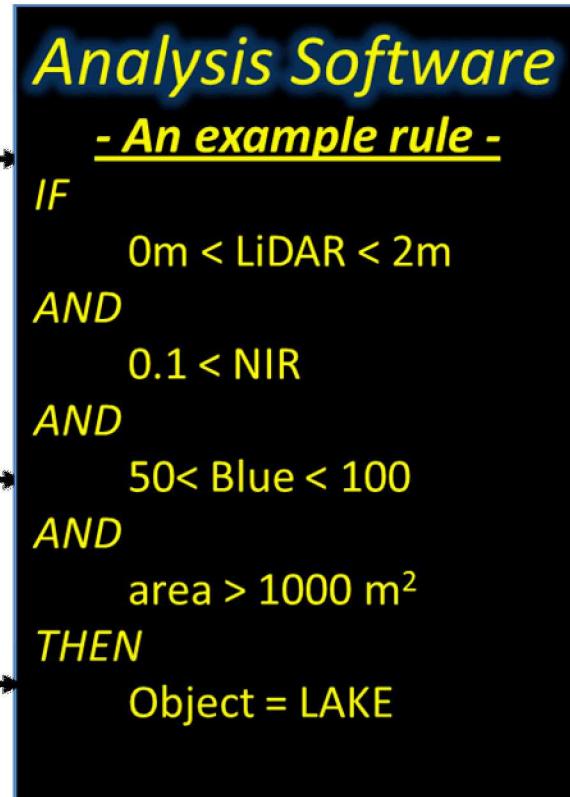


Region Classification

We use eCognition* software to create land classification maps:



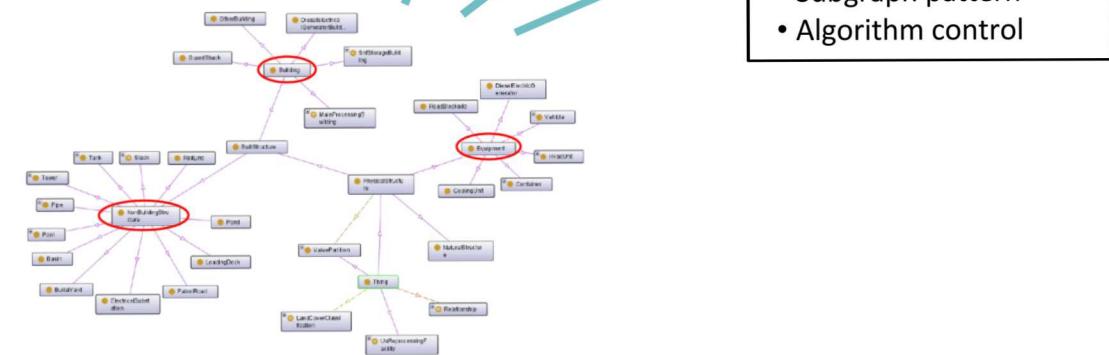
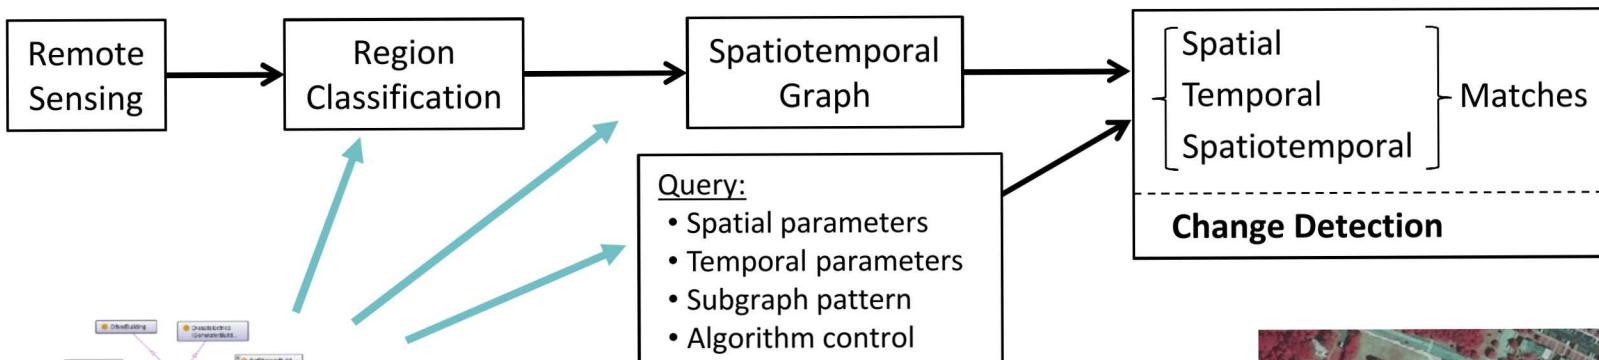
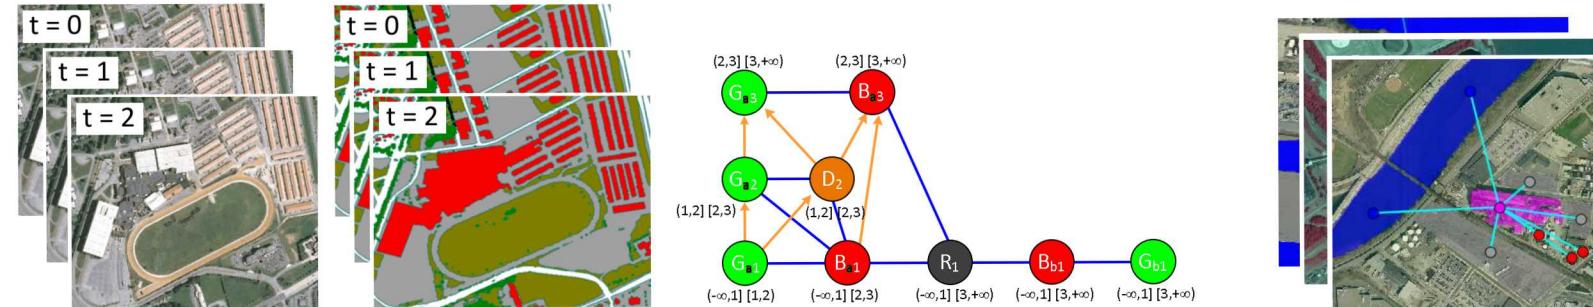
Rule Sets



Posterized image

*eCognition is used throughout the project, by LLNL, SNL and U. Vermont

Geospatial-Temporal Graph Data Flow



Ontology

Input Data

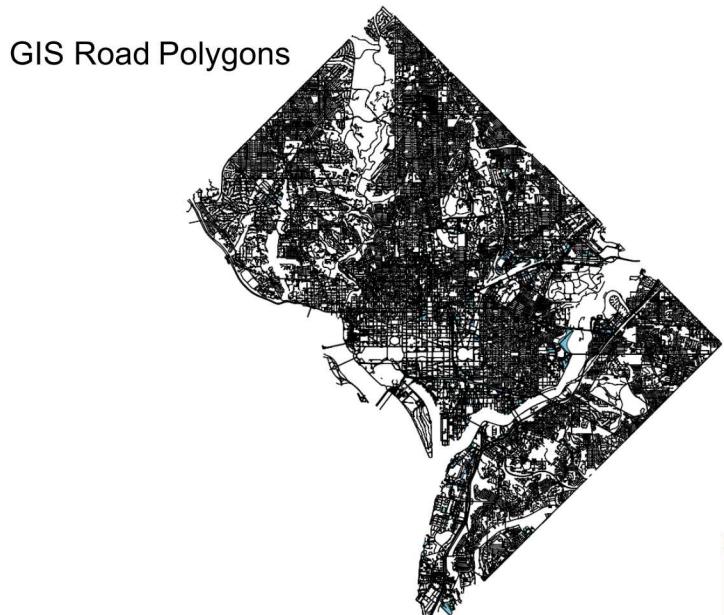
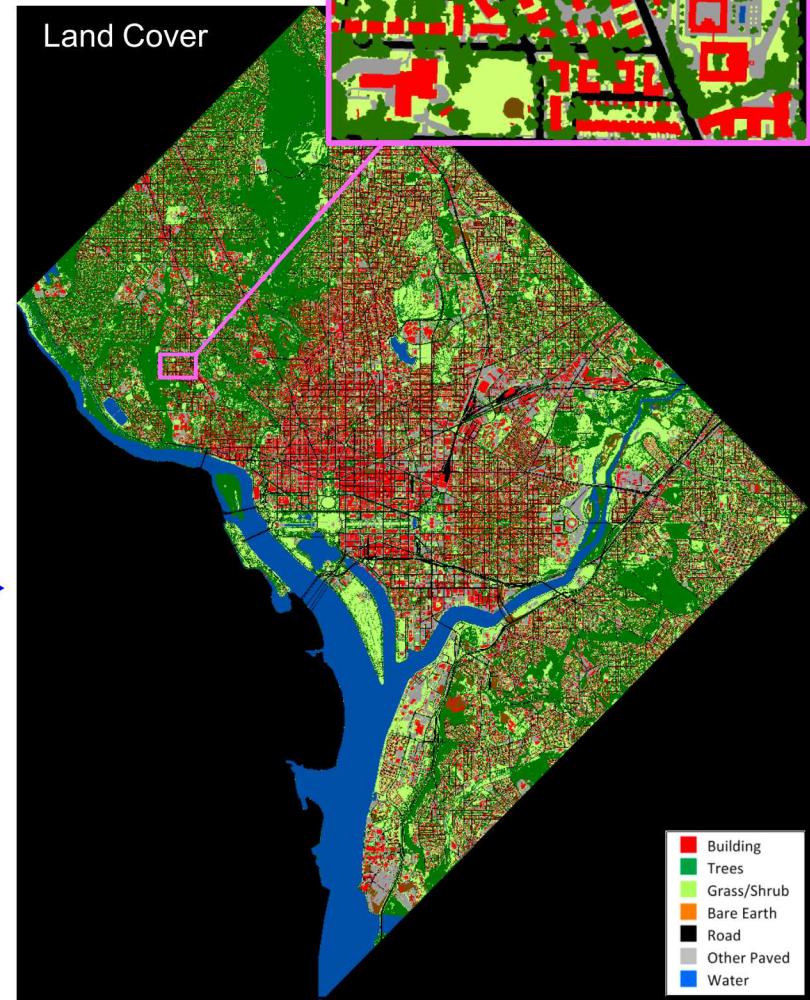
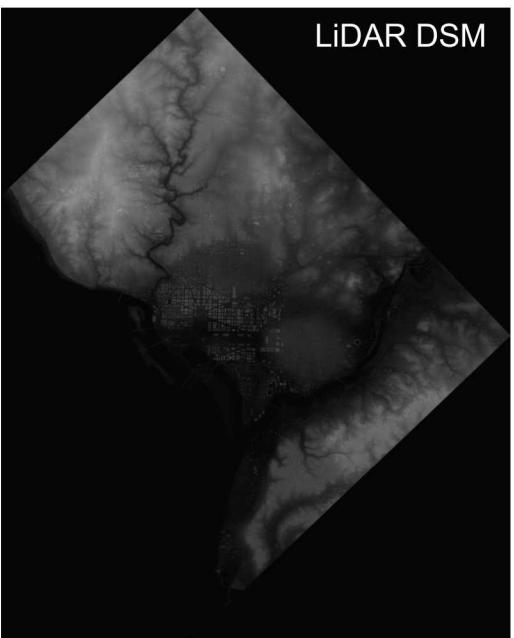


Image data and processing by University of Vermont
Spatial Analysis Lab [O'Neil-Dunne 2012].

Systematic Land Cover Classification

Steps:

1. Assemble input:

- RGB+IR imagery, near nadir.
- LiDAR data.
- GIS road centerlines or polygons (if available).

2. Construct normalized digital surface model (nDSM).

3. Using nDSM, split image into “Tall” vs. “Short.”

4. Using RGB+IR, split out “Tree” from Tall.

5. Using geometric morphology and LiDAR texture, split out “Building” from Tall.

6. Using morphology and LiDAR, split remaining Tall into “Wall” and “Other Structure.”

7. Using GIS road data, split out “Road” from Short.

8. Using RGB+IR, split remaining Short into “Grass/Shrub,” “Dirt,” “Water,” and “Other Paved.”

9. Output posterized file.

Work is proceeding at UVM, LLNL, and SNL evaluating this procedure.

Example from UVM

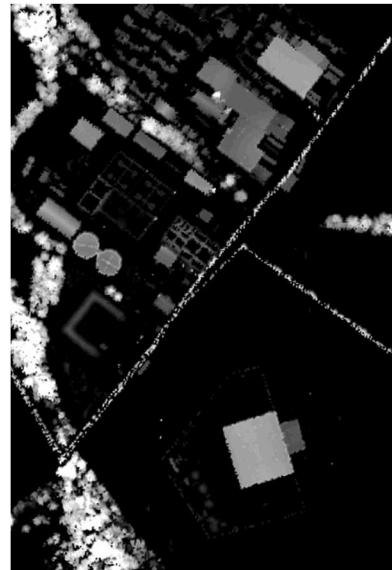
1. Image (True Color)



2. Image (Color Infrared)



3. LiDAR nDSM



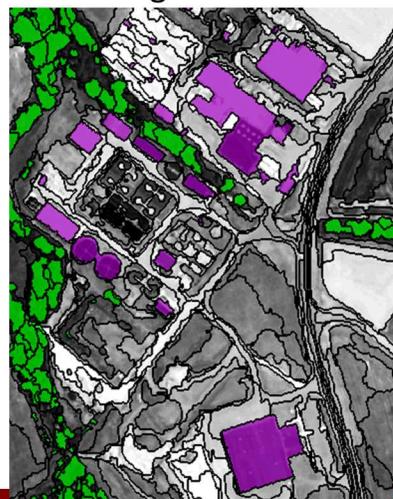
4. Slope from nDSM



5. Tall



6. Buildings and Trees



7. Grass and Shrubs

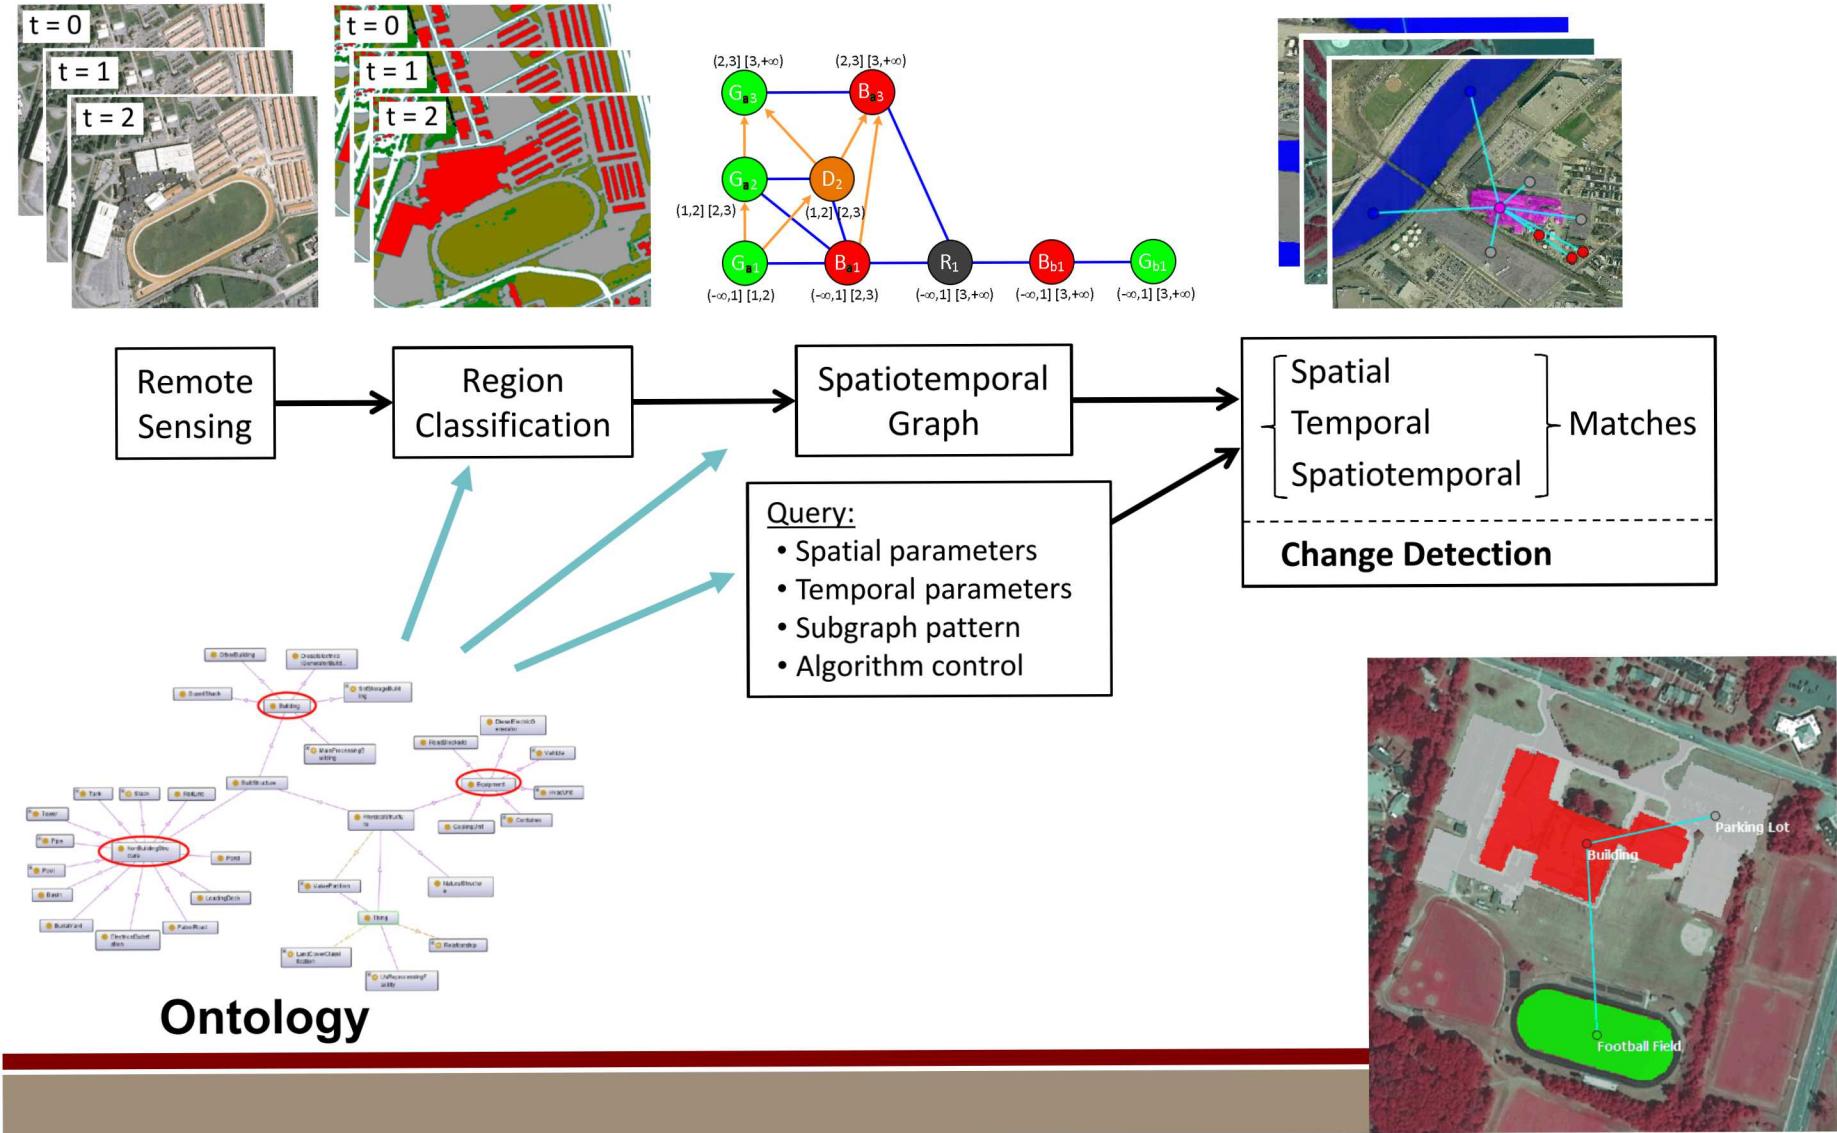


8. Final Land Cover



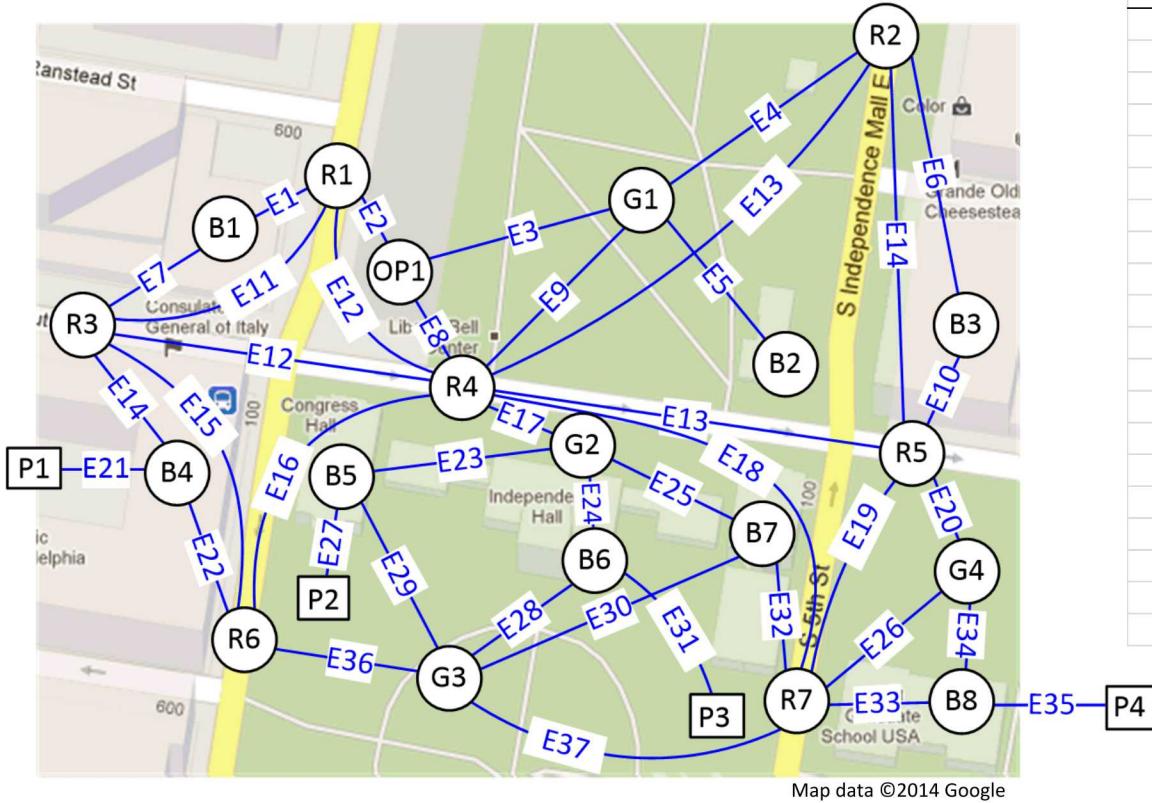
- Bare Soil
- Buildings
- Grass/Shrub
- Other Paved
- Railroads
- Roads
- Tree Canopy
- Water

Geospatial-Temporal Graph Data Flow



Example Geospatial Semantic Graph

Independence Hall, Philadelphia:



Point node table:

id	Name	Address	Latitude	Longitude
P1	Consulate of Italy	150 S. Independent Mall West #1026	-75.14895	39.94884
P2	Congress Hall	41 N 6th Street	-75.14920	39.94899
P3	Independence Hall	520 Chestnut Street	-75.15000	39.94889
P4	Graduate School USA	150 S. Independence Mall West #674	-75.15090	39.94819

Region node table:

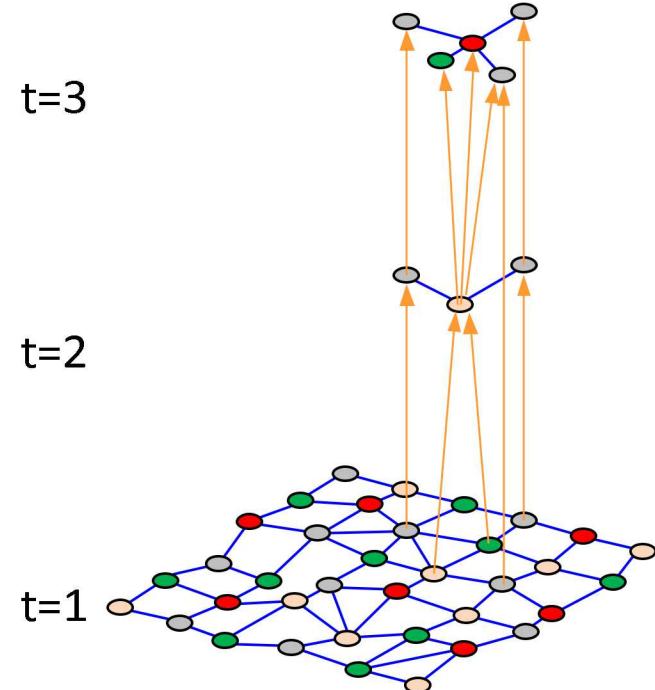
id	type	area	centroid x	centroid y
B1	building	3200	-75.14900	39.94939
R1	road	1800	-75.14910	39.94949
OP1	paved	4700	-75.14935	39.94934
G1	grass	22000	-75.15010	39.94944
R2	road	1900	-75.15060	39.94999
R3	road	1100	-75.14885	39.94934
R4	road	2200	-75.14980	39.94924
B2	building	780	-75.15045	39.94931
B3	building	6000	-75.15075	39.94944
B4	building	12000	-75.14895	39.94884
B5	building	2100	-75.14920	39.94899
G2	grass	7700	-75.14990	39.94906
R5	road	870	-75.15065	39.94896
B6	building	2000	-75.15000	39.94889
B7	building	3150	-75.15040	39.94884
G4	grass	15300	-75.15080	39.94869
R6	road	1970	-75.14905	39.94844
G3	grass	25000	-75.14960	39.94829
R7	road	1810	-75.15050	39.94834
B8	building	2700	-75.15090	39.94819

Edge table:

edge_id	node_1	node_2
E1	B1	R1
E2	R1	OP1
E3	OP1	G1
E4	G1	R2
E5	G1	B2
E6	R2	B3
E7	R3	B1
E8	OP1	R4
E9	R4	G1
.	.	.
.	.	.
.	.	.

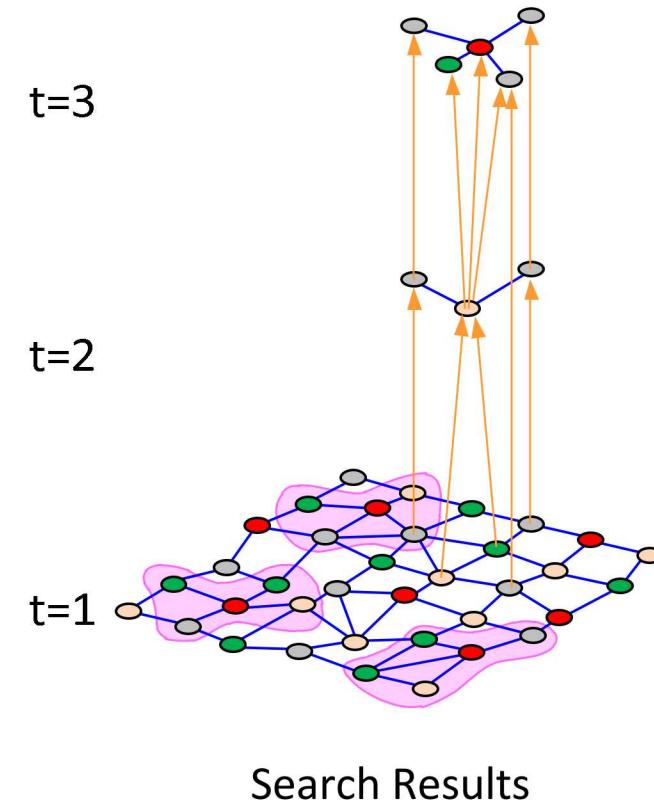
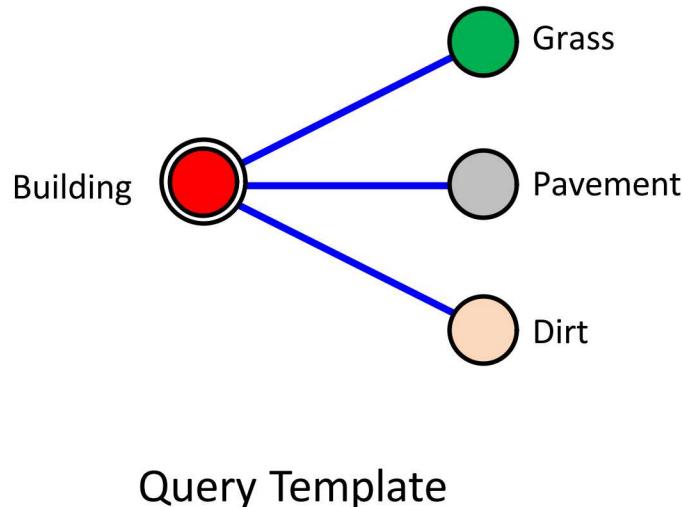
Representing Change Over Time

- Encode change:
 - Node attributes include duration seen.
 - Only construct new nodes for changes.
 - “Changed-to” arcs encode time evolution.
 - Graph complexity focuses on change areas.



Signature Search

- A signature encodes a desired question.
- Example: “Where are buildings with nearby grass, pavement, and dirt?



Overview

- Motivation.
- Computation.
- ■ Example Results.
- Extensions.
- Discussion.

Overview

- Motivation.
- Computation.
- Example Results.

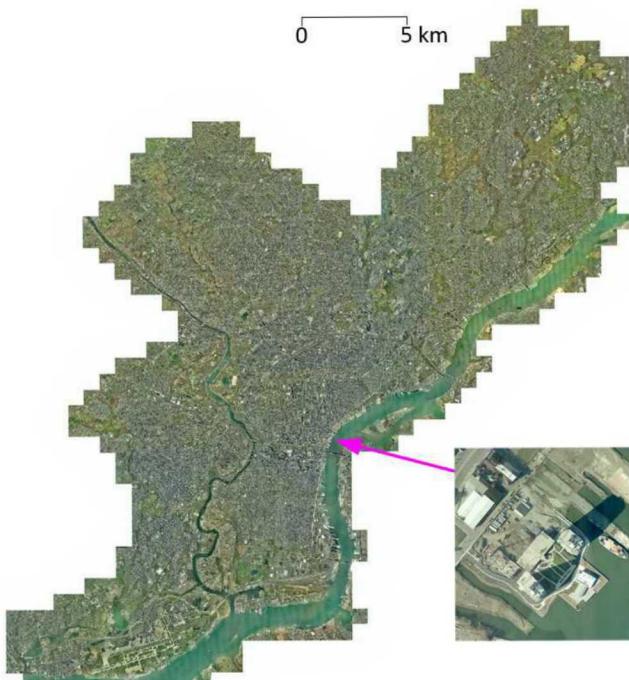
→ ■ Data.

- Power plants.
- Refineries.
- Change analysis.
- New complex.

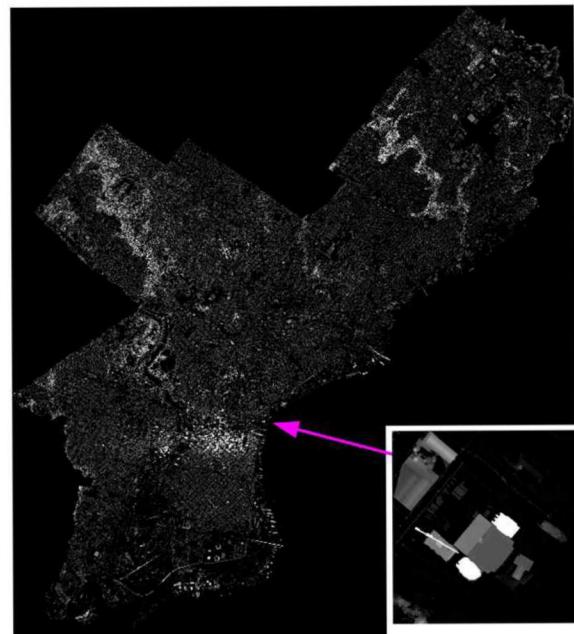
- Extensions.
- Discussion.

Philadelphia 2008

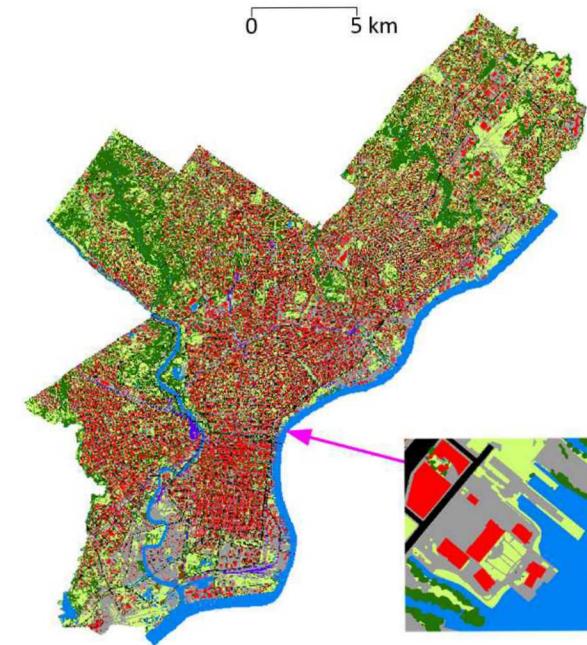
Primary input:



Optical Image



LiDAR nDSM



Land Cover

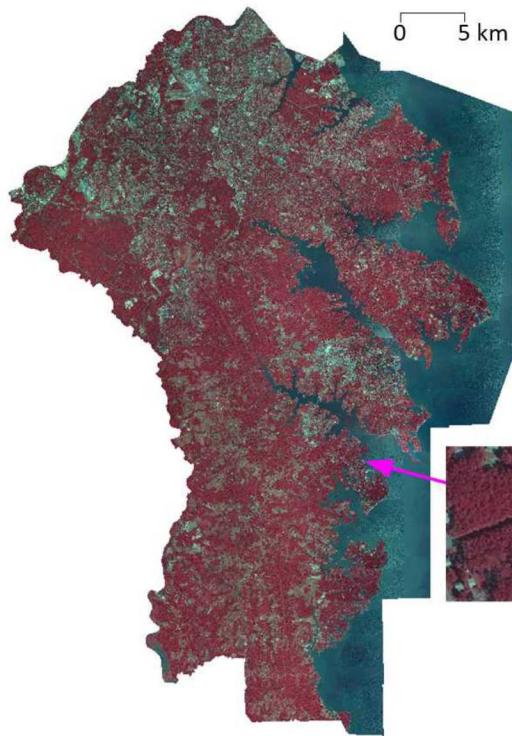
Pixel size 0.1 m
307,531 × 330,033 pixels
(101.5 Gpix)
7,669 MB

Pixel size 0.3 m
89,540 × 100,294 pixels
(9.0 Gpix)
2,084 MB

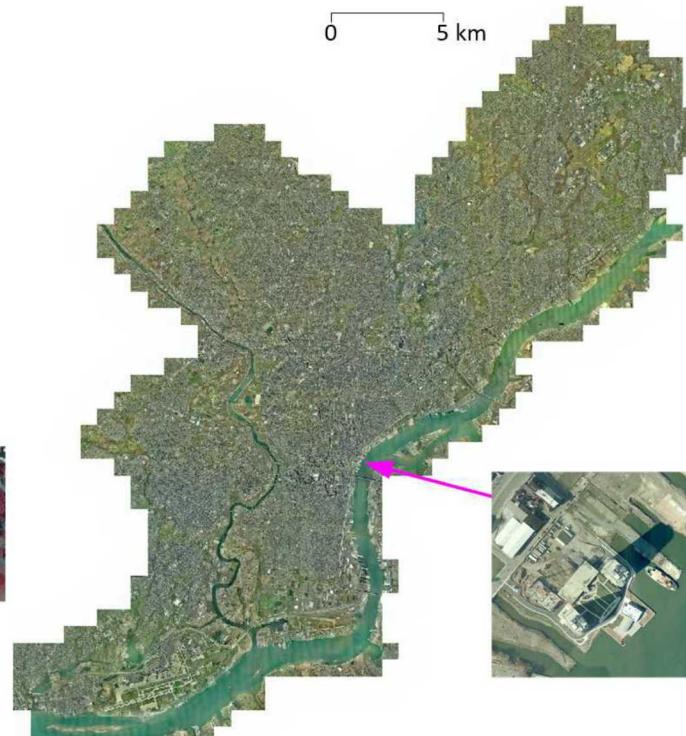
Pixel size 0.3 m
89,548 × 100,303 pixels
(9.0 Gpix)
8,775 MB

Three Data Regions

Search results:



Anne Arundel County, MD



Philadelphia, PA



Washington, DC

Total:

2,067 km² total area
135 billion Pixels
3.6 million Features

Total file size was about 88 GB.

Overview

- Motivation.
- Computation.
- Example Results.
 - Data.
- ■ Power plants.
 - Refineries.
 - Change analysis.
 - New complex.
- Extensions.
- Discussion.

Power Plant Search

Query specification:

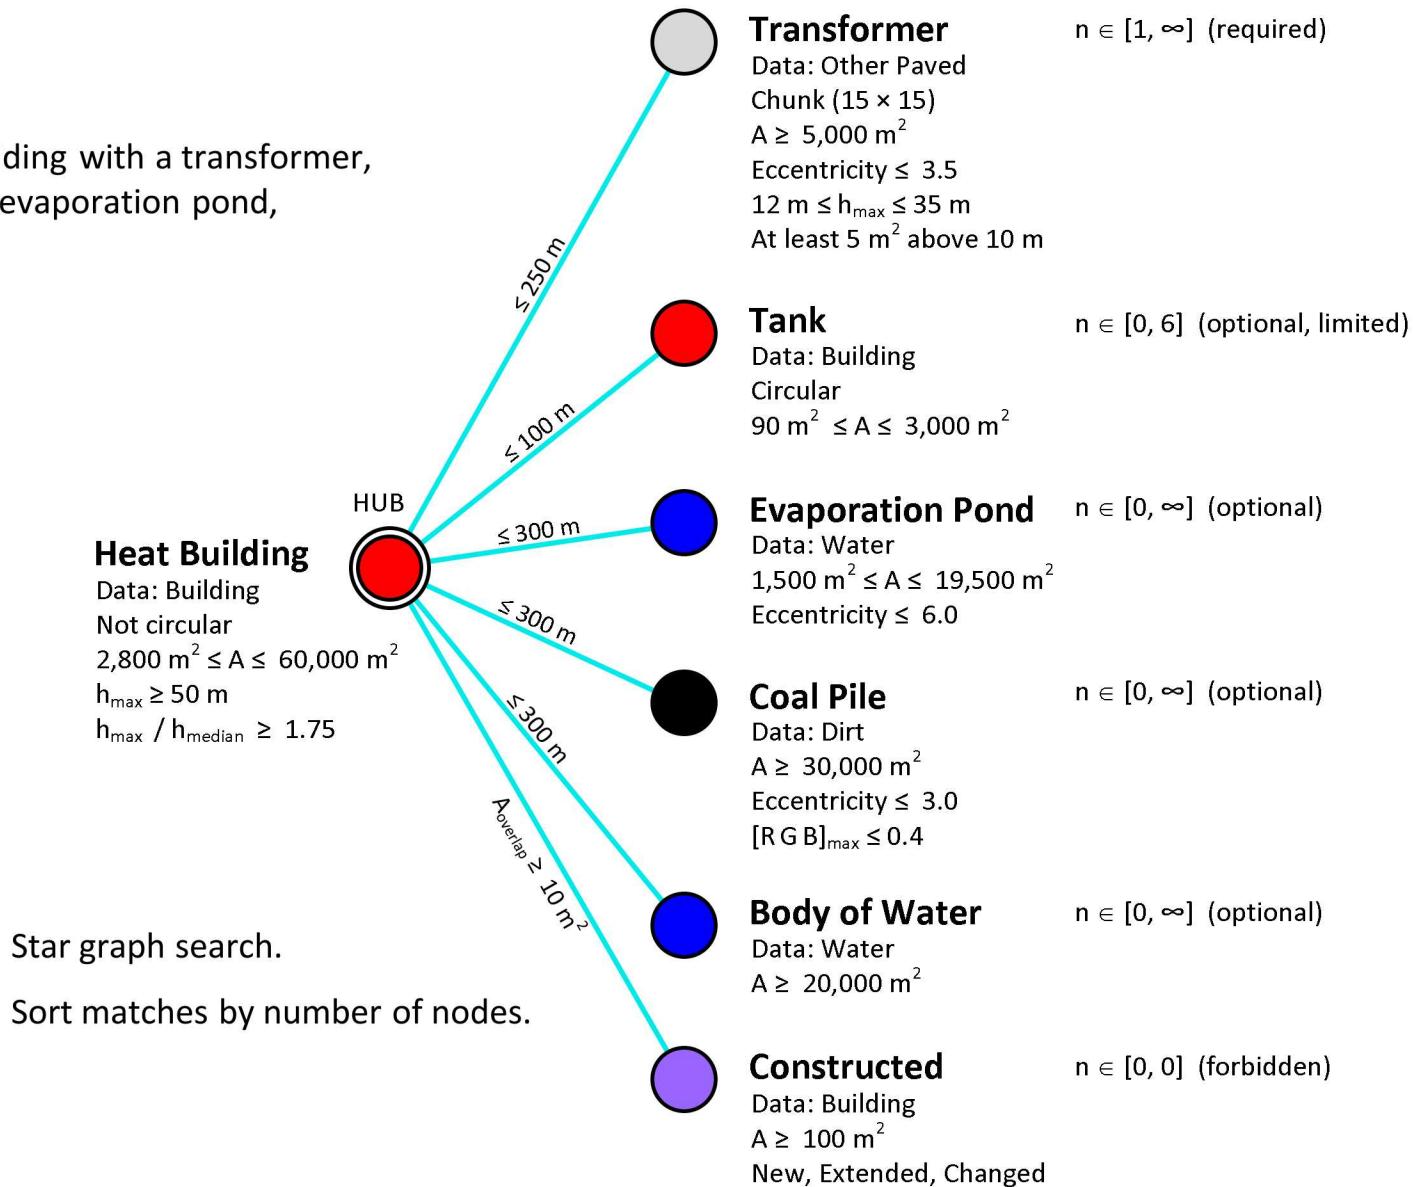
A power plant is a heat building with a transformer, and optional storage tank, evaporation pond, coal pile, body of water.

Question semantics:

Heat Building
 Transformer
 Cooling Tower
 Evaporation Pond
 Body of Water
 Coal Pile
 Storage Tanks
 Processing Tower
 Pipe Network
 :
 :

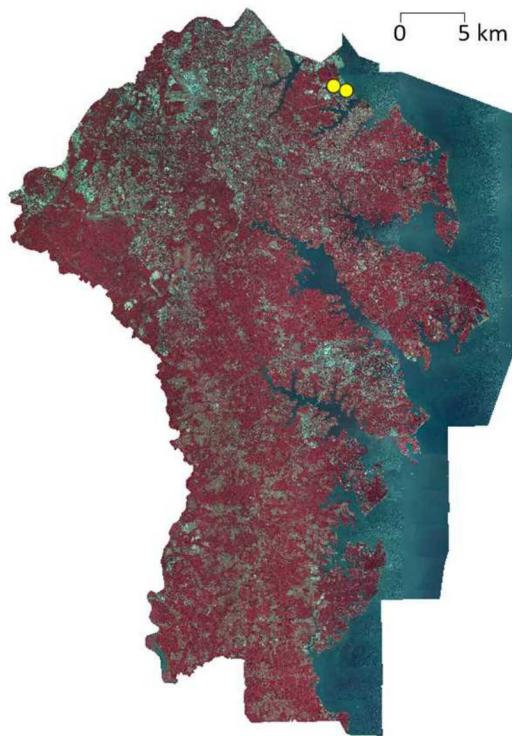
Data semantics:

Building
 Road
 Other Paved
 Grass/Shrub
 Trees
 Dirt
 Water

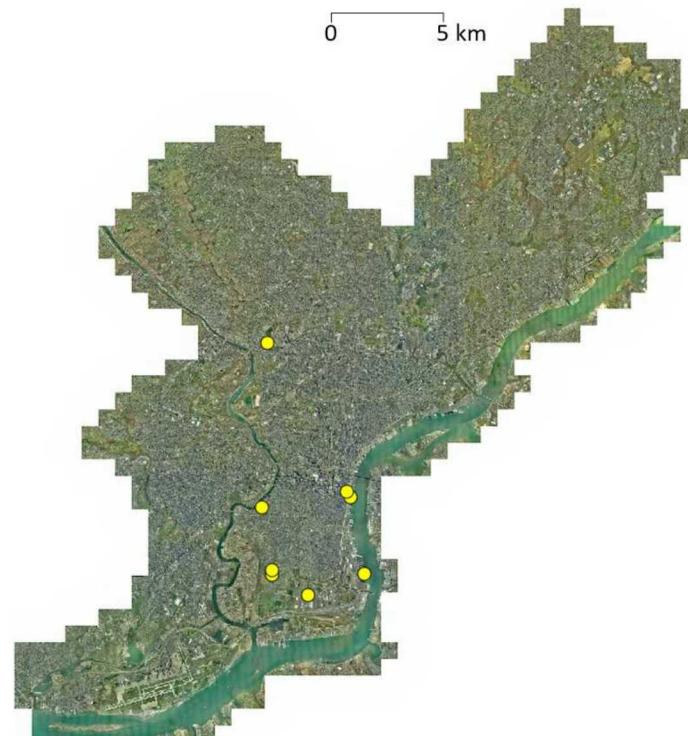


Power Plant Search

Search results:



Anne Arundel County, MD



Philadelphia, PA



Washington, DC

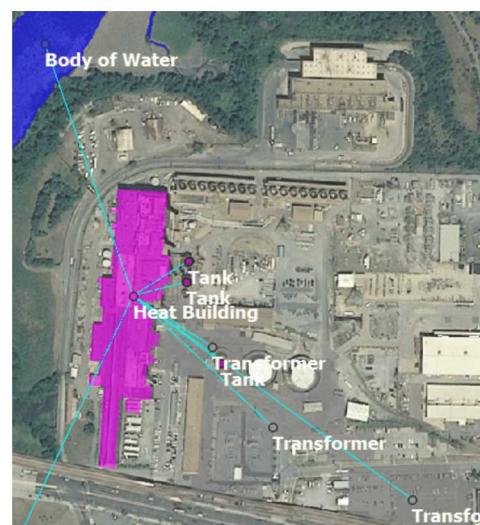
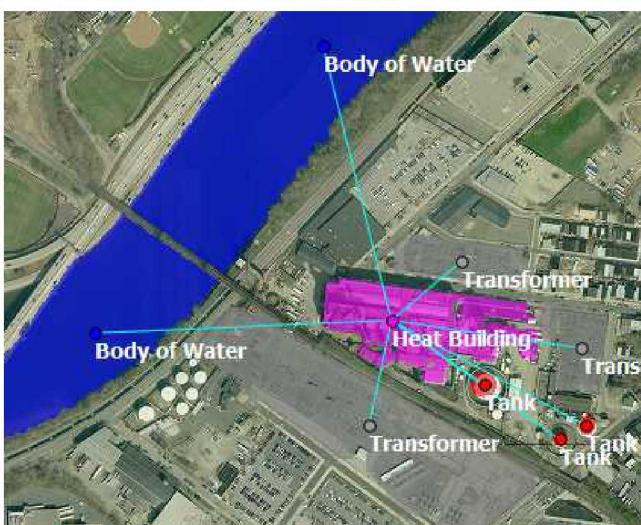
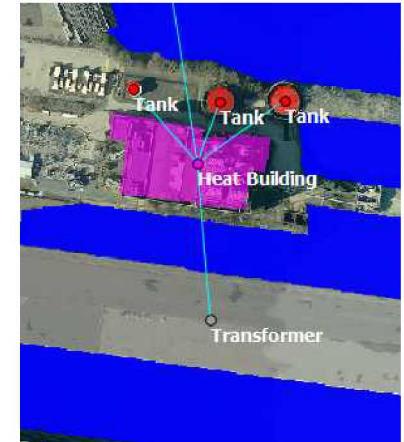
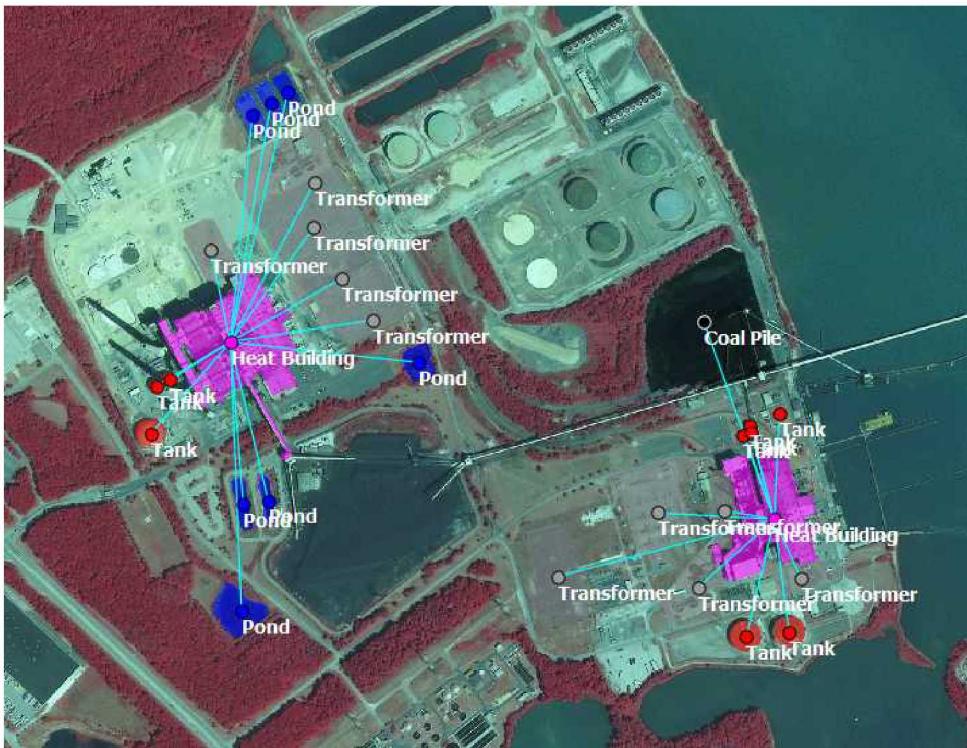
Input:

2,067 km² total area
135 billion Pixels
3.6 million Features

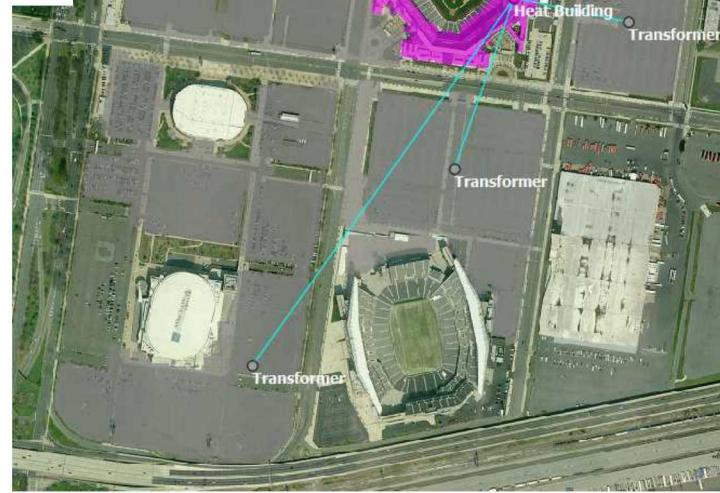
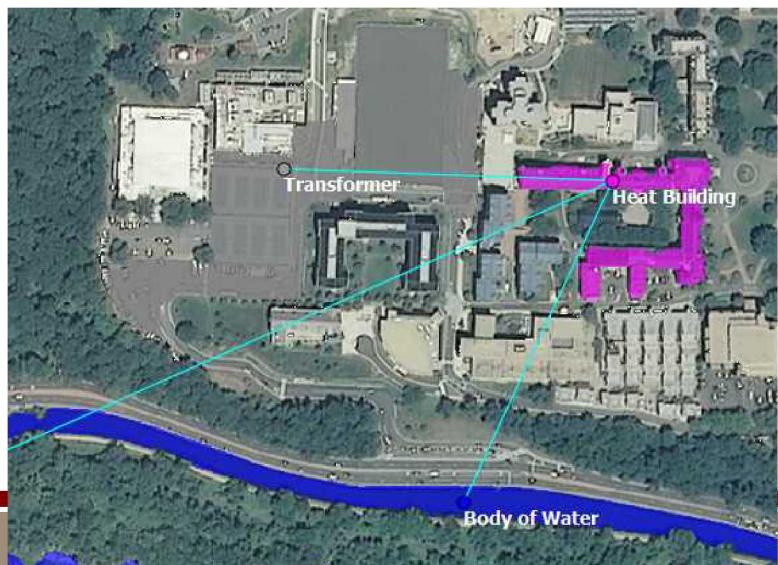
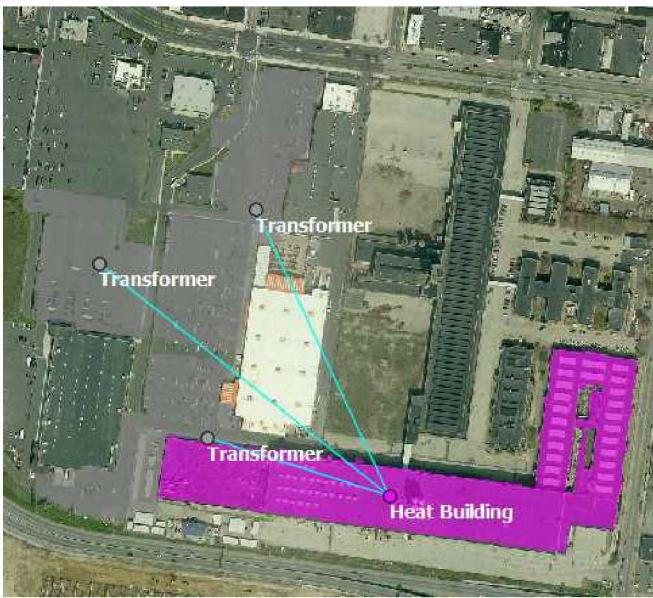
Output:

6 True positives
9 False positives
2 False negatives
(2 Invisible)

Power Plant Results: True Positives



Power Plant Results: False Positives



A better transformer filter would eliminate these.

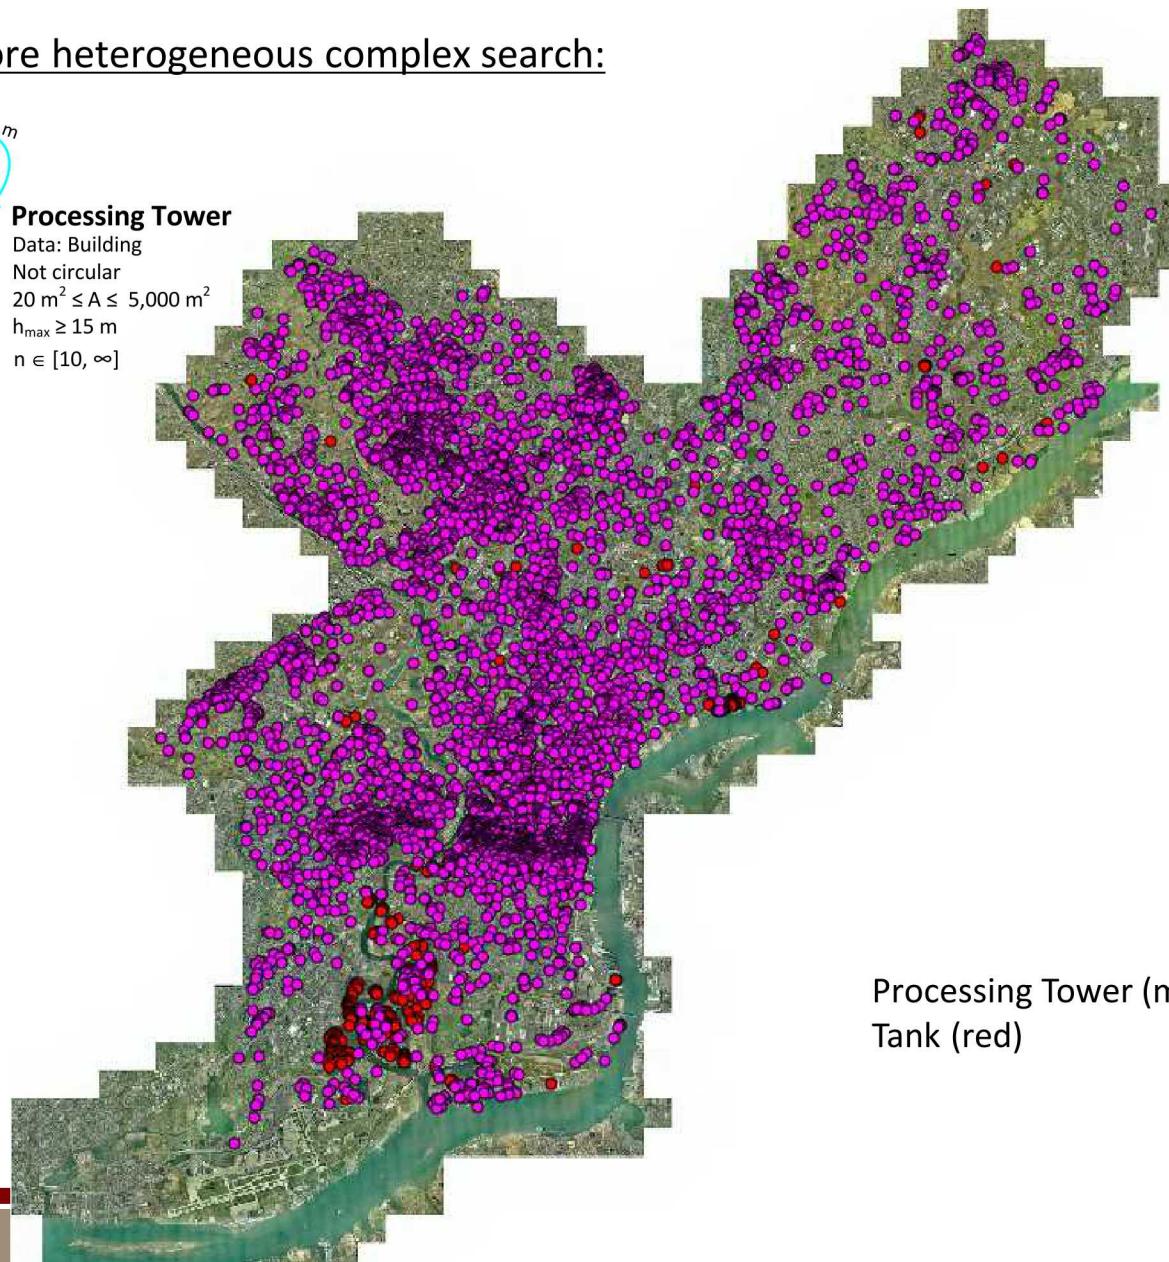
Overview

- Motivation.
- Computation.
- Example Results.
 - Data.
 - Power plants.
 - Refineries.
 - Change analysis.
 - New complex.
- Extensions.
- Discussion.



Large Refinery Search

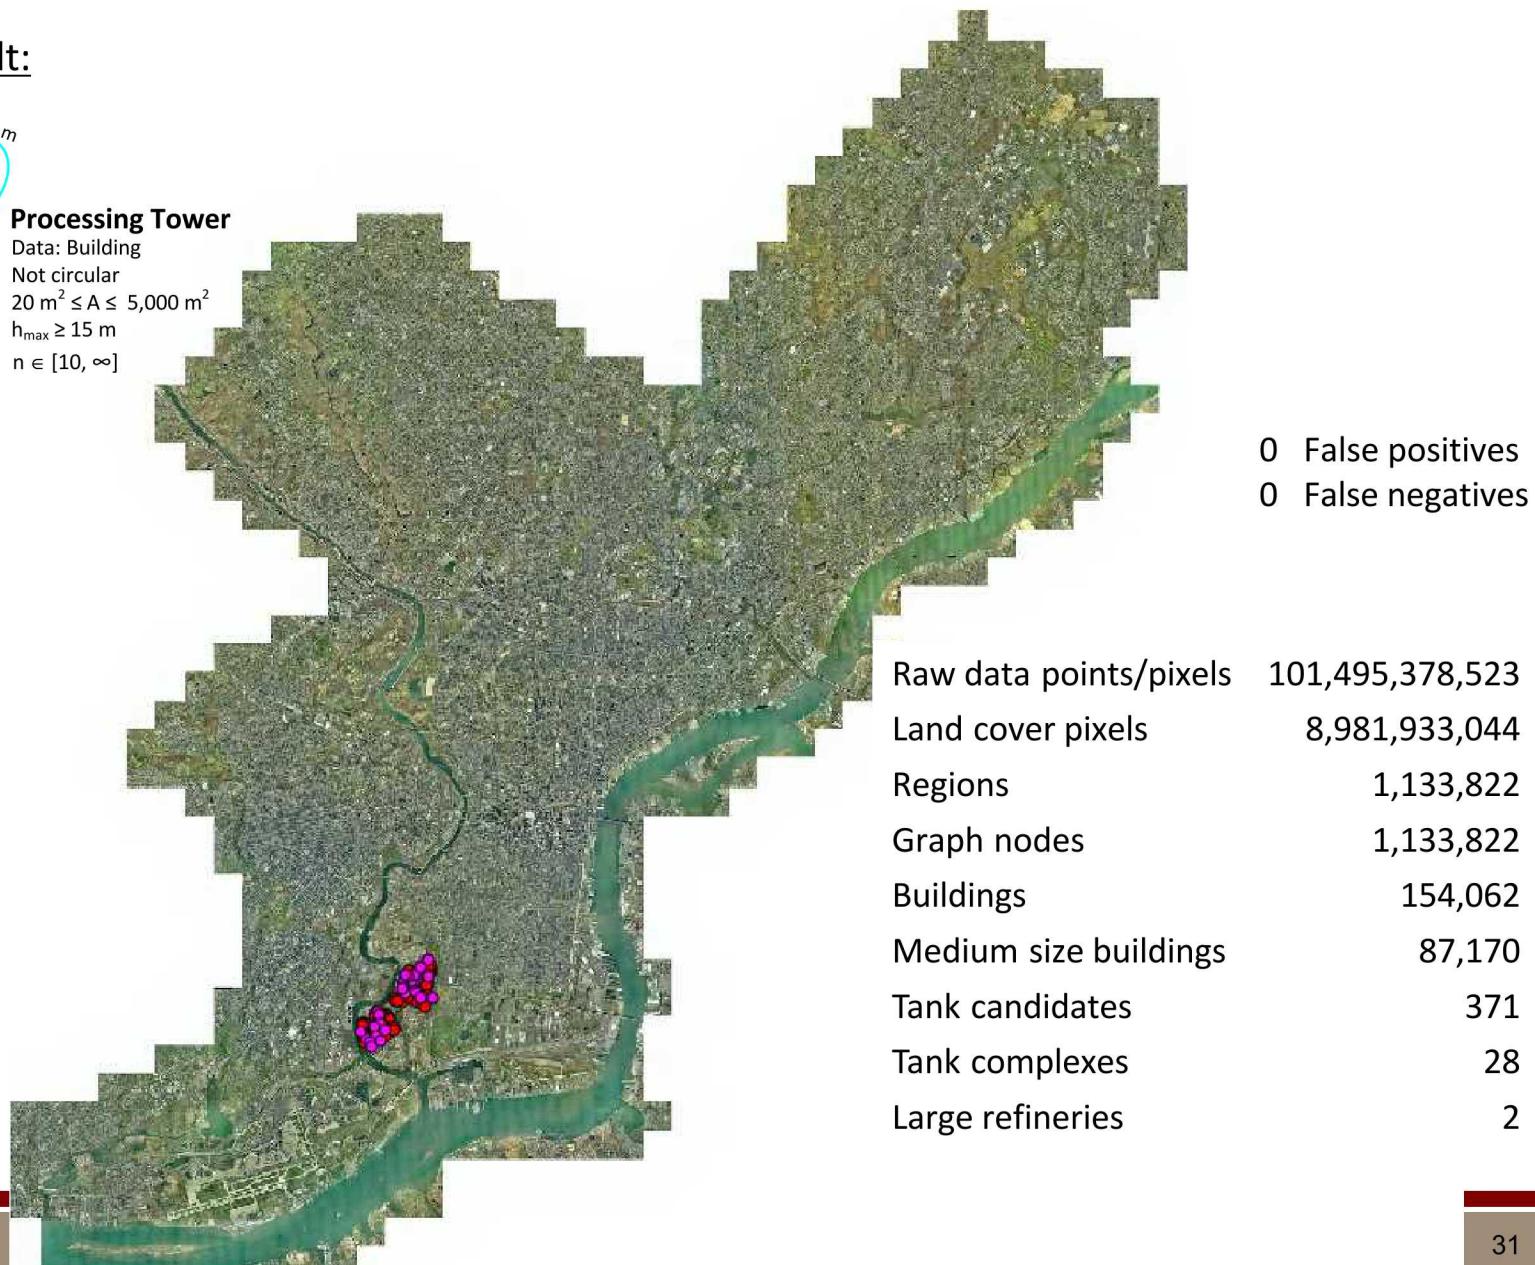
SearchGraph, before heterogeneous complex search:



Processing Tower (magenta) 4,909
Tank (red) 371

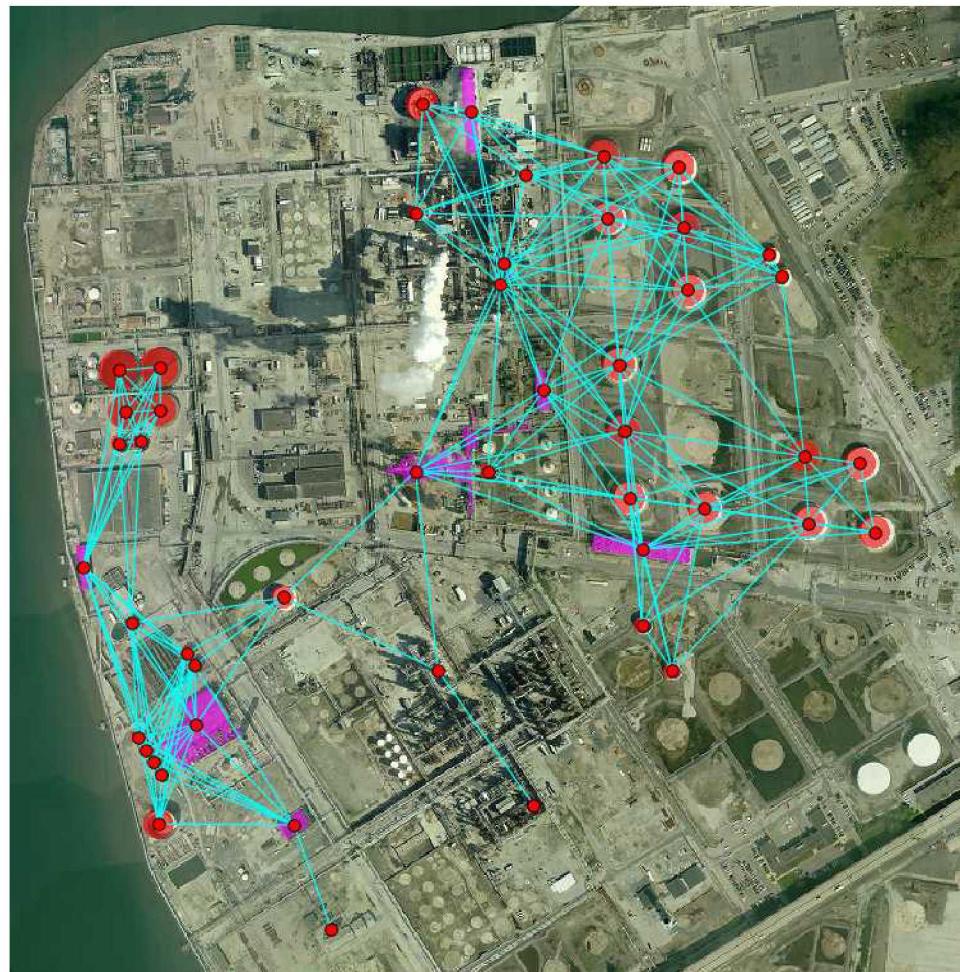
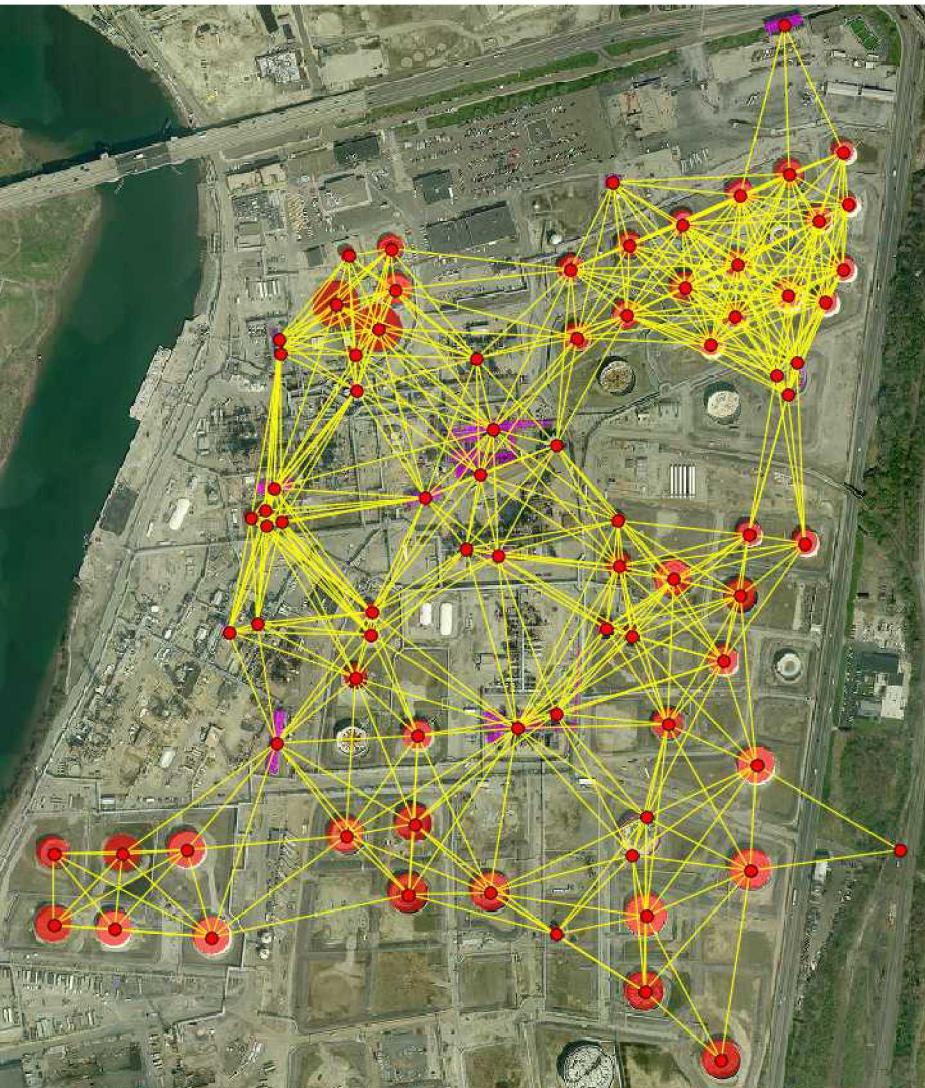
Large Refinery Search

Graph search result:



Large Refinery Search

Refineries found:

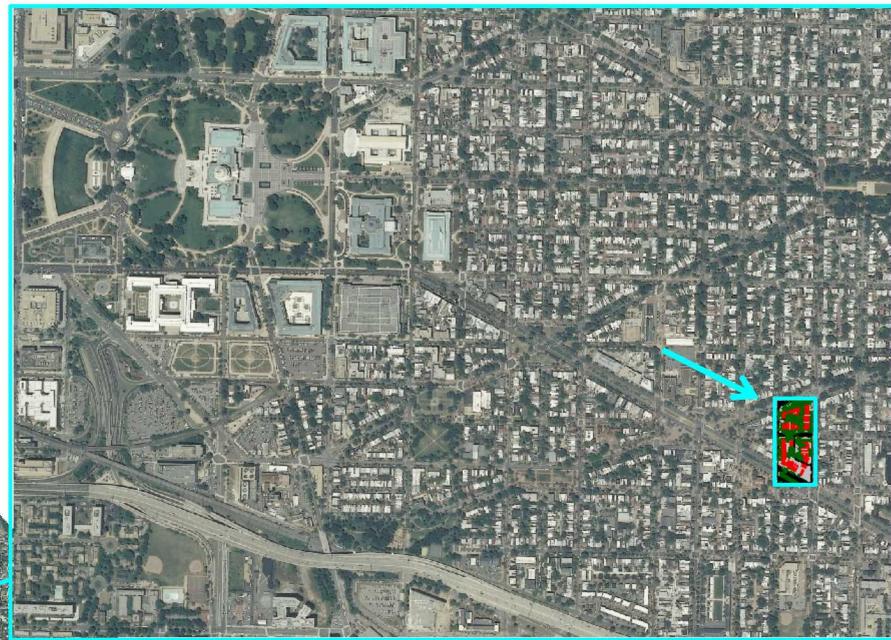


Overview

- Motivation.
- Computation.
- Example Results.
 - Data.
 - Power plants.
 - Refineries.
- ■ Change analysis.
- New complex.
- Extensions.
- Discussion.

Simple Change Example

Context:



Simple Change Example

(NA-22 Project)

Resulting land cover:



2006

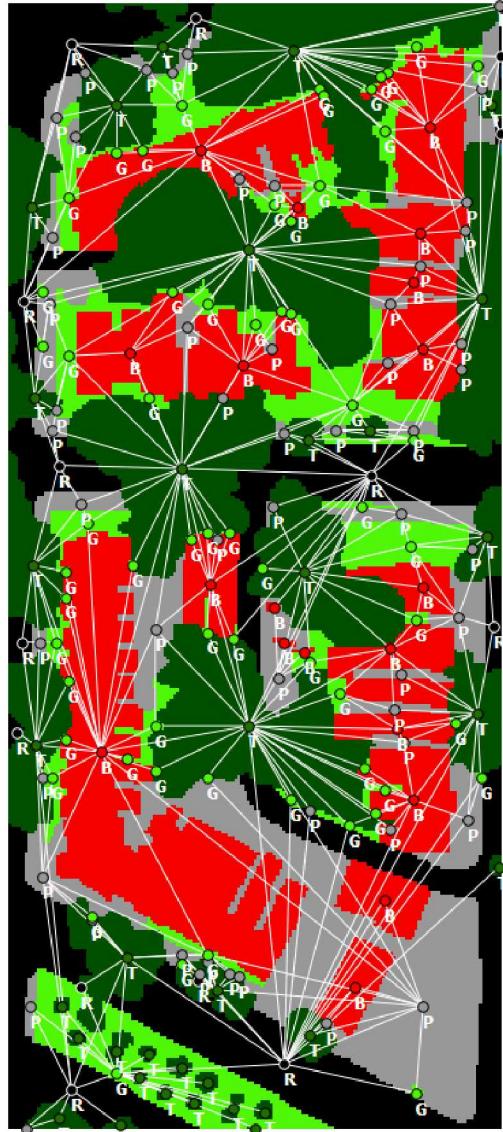


2011

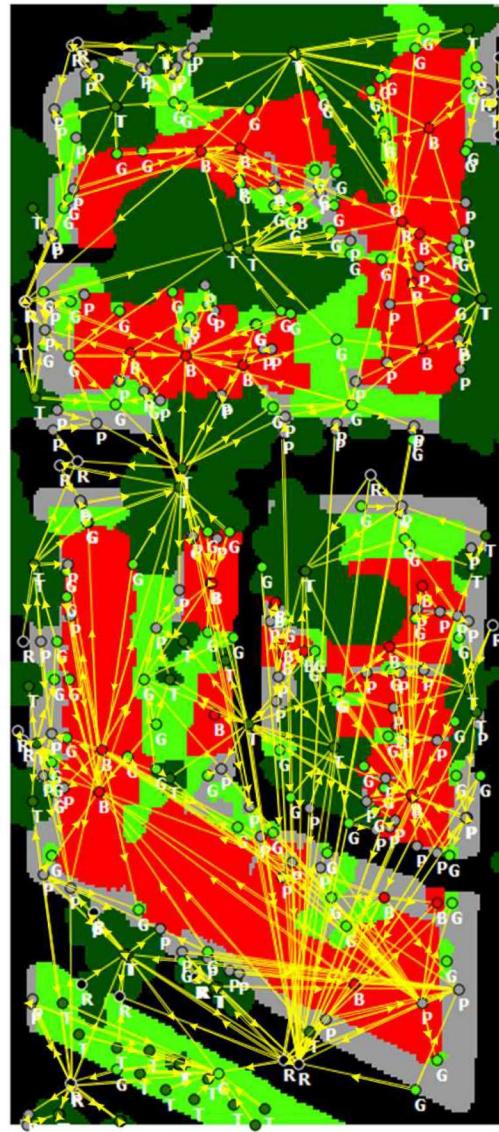
StoredGraph in Three Time Slices

(NA-22 Project)

Geospatial-temporal graph:



2006 Adjacency



2006 → 2011 Change



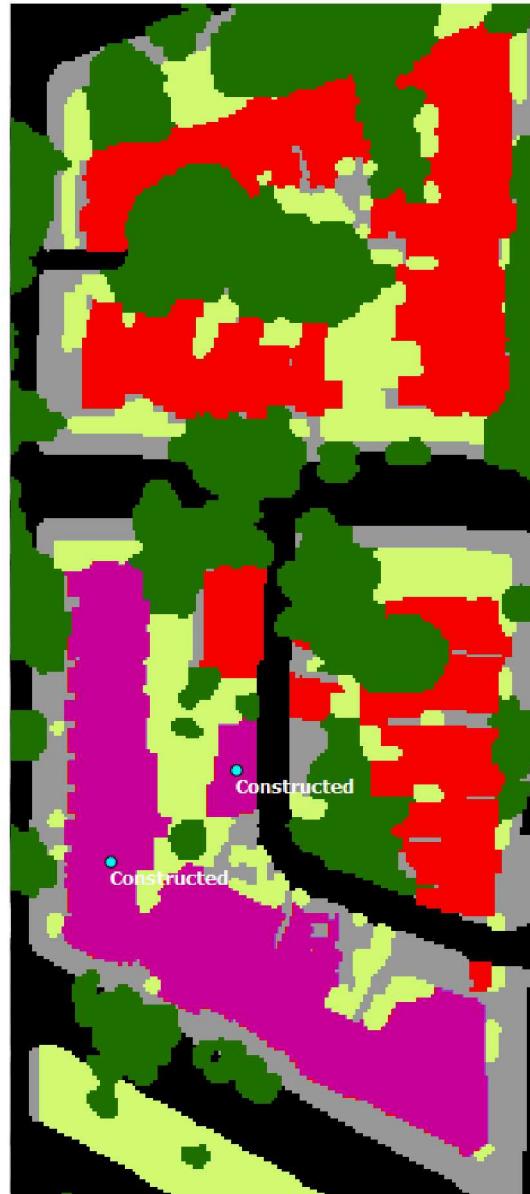
2011 Adjacency

342 nodes
808 adjacency edges
559 change edges
0 distance edges

Differentiating Important Change

(NA-22 Project)

Graph-based change analysis results:



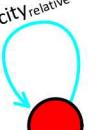
Overview

- Motivation.
- Computation.
- Example Results.
 - Data.
 - Power plants.
 - Refineries.
 - Change analysis.
 - New complex.
- ■ Extensions.
- Discussion.

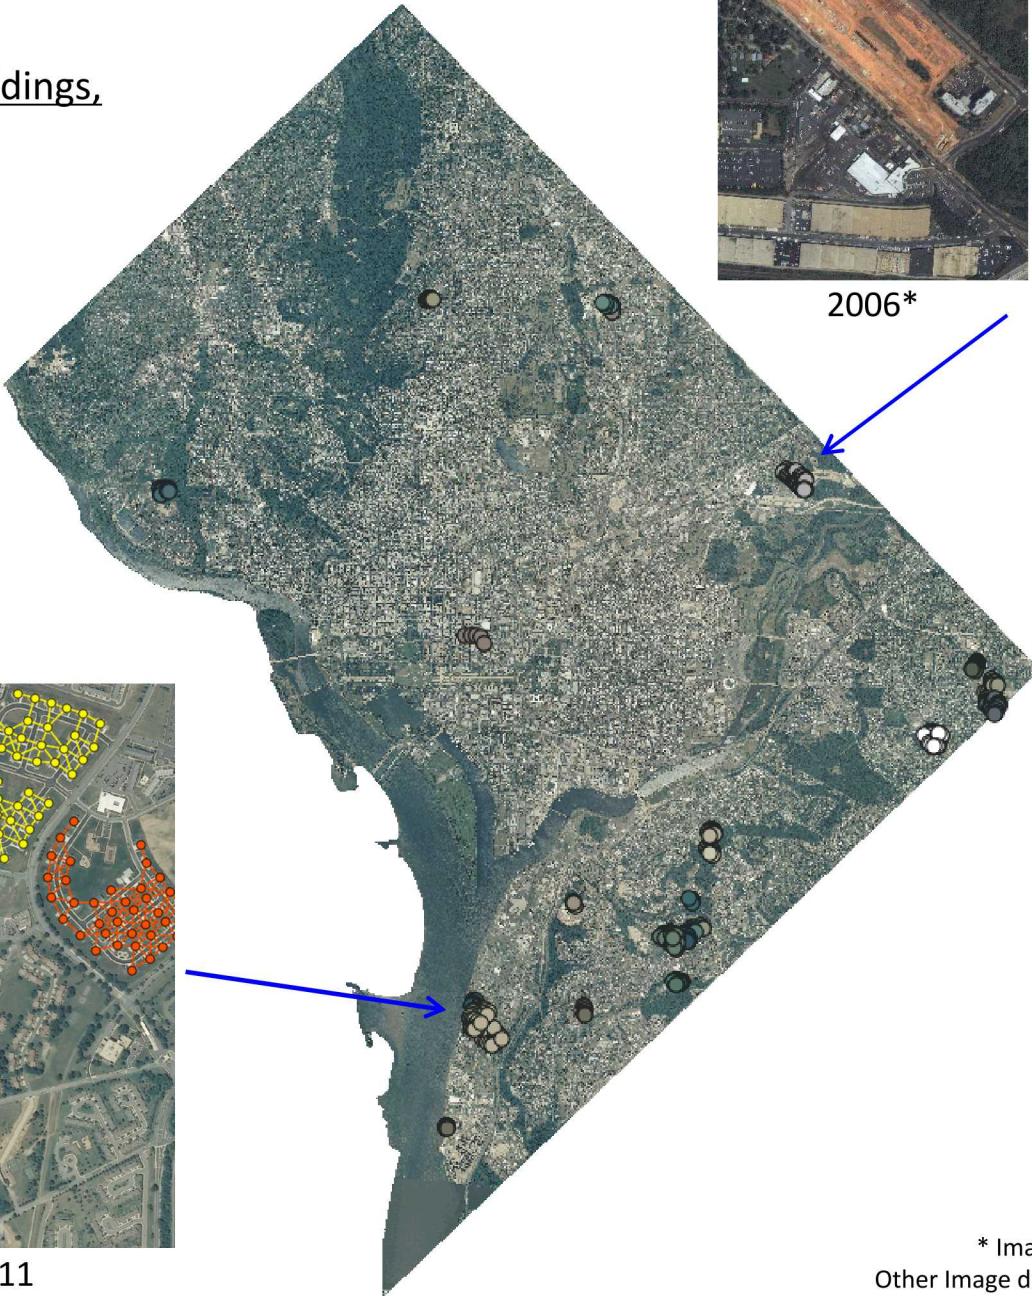
New Complexes

Seek complexes of new buildings,
across the entire city:

$\leq 40\text{ m}$
 $A_{\text{relative}} \leq 1.5\times$
 $\text{Eccentricity}_{\text{relative}} \leq 1.5\times$



Constructed
Data: Building
Exists now
 $A \geq 100\text{ m}^2$
New, Extended, Changed
 $n \in [5, \infty]$



2006*



2011



2006*



2011

* Image from DigitalGlobe.
Other Image data provided by UVM.

Overview

- Motivation.
- Computation.
- Example Results.
- ▪ Extensions.
- Discussion.

Overview

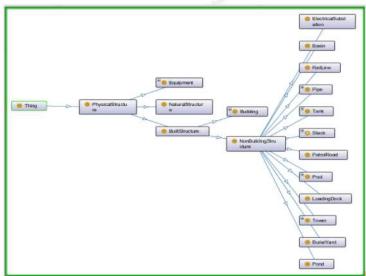
- Motivation.
- Computation.
- Example Results.
- Extensions.
- ▪ Query generation.
 - Semantic refinement: Chunks, path network.
 - Search writeback and re-use.
 - Quality scoring.
 - Activity analysis.
- Discussion.

Approaches to Query Construction

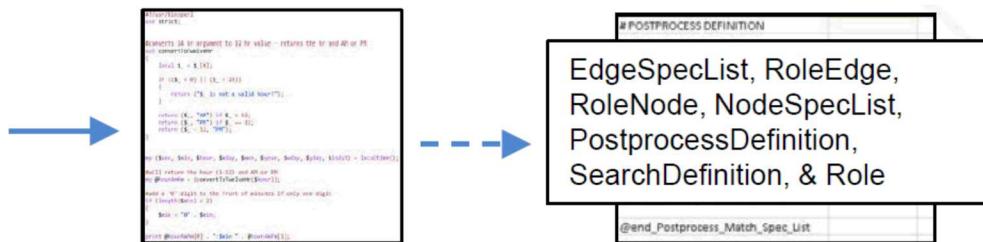
Methods:

- GeoQuestion csv file code.
- Definition with GUI.
- Query-by-example.
- Generation from ontology.

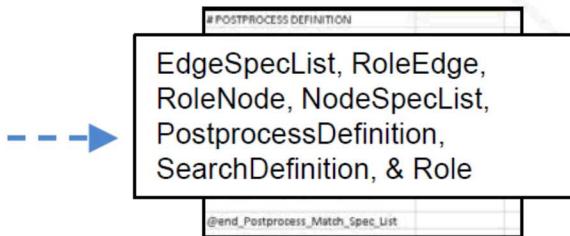
Ontology to Query Generation



Knowledge Base
Protégé
Craig Blackhart



O2Geo
Perl
Shelly Spearing



GeoQuestion input
(7 CSVs / query)
MS Excel
SNL and LANL

Geospatial Data & posterizations
LLNL & U.Vt

Cyclic Process
Discrete Releases
(e.g., “v1.0.4”)



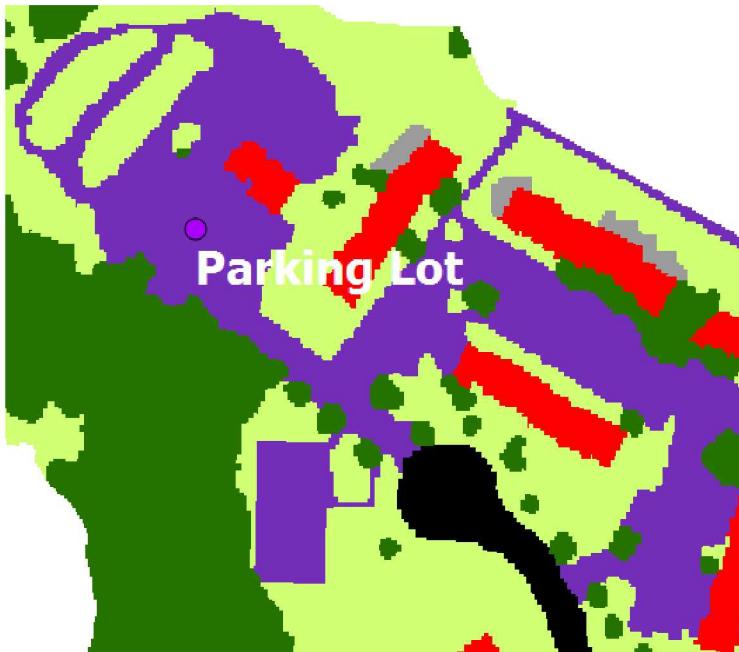
GeoGraphy
Randy Brost, et al.

Search output
List of matches

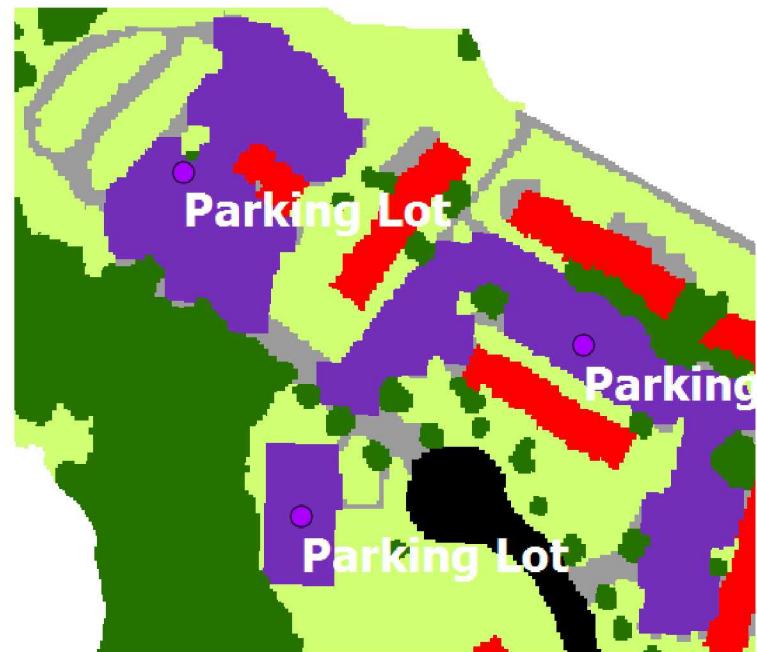
Chunk Operation

Chunking breaks up linked regions into significant sub-regions:

Before Chunking



After Chunking



Implemented by a geometric shrink/grow operation, with controllable magnitude.

Multi-Step GeoSpatial-Temporal Search

Pre-processing:



Interactive:



Search 1



Search 2



Search 3



Search 4

(Using results from 1, 2, 3)

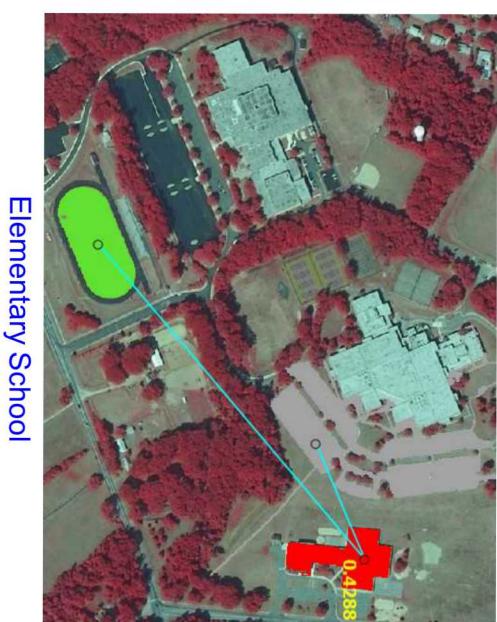
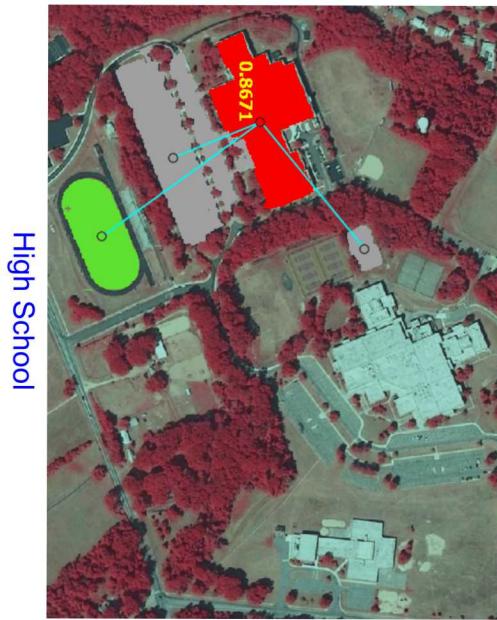
Search results can be written back to the graph, used as components of later searches.

Advantages:

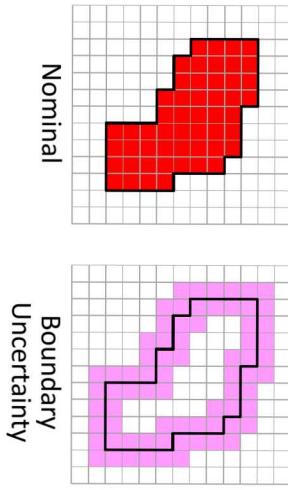
- Batch “micro-searches.”
- Smaller, modular queries.
- Log analyst feedback.
- Search re-use.
- Hierarchical semantics.
- Expands scope of possible searches.

Match Quality Scoring

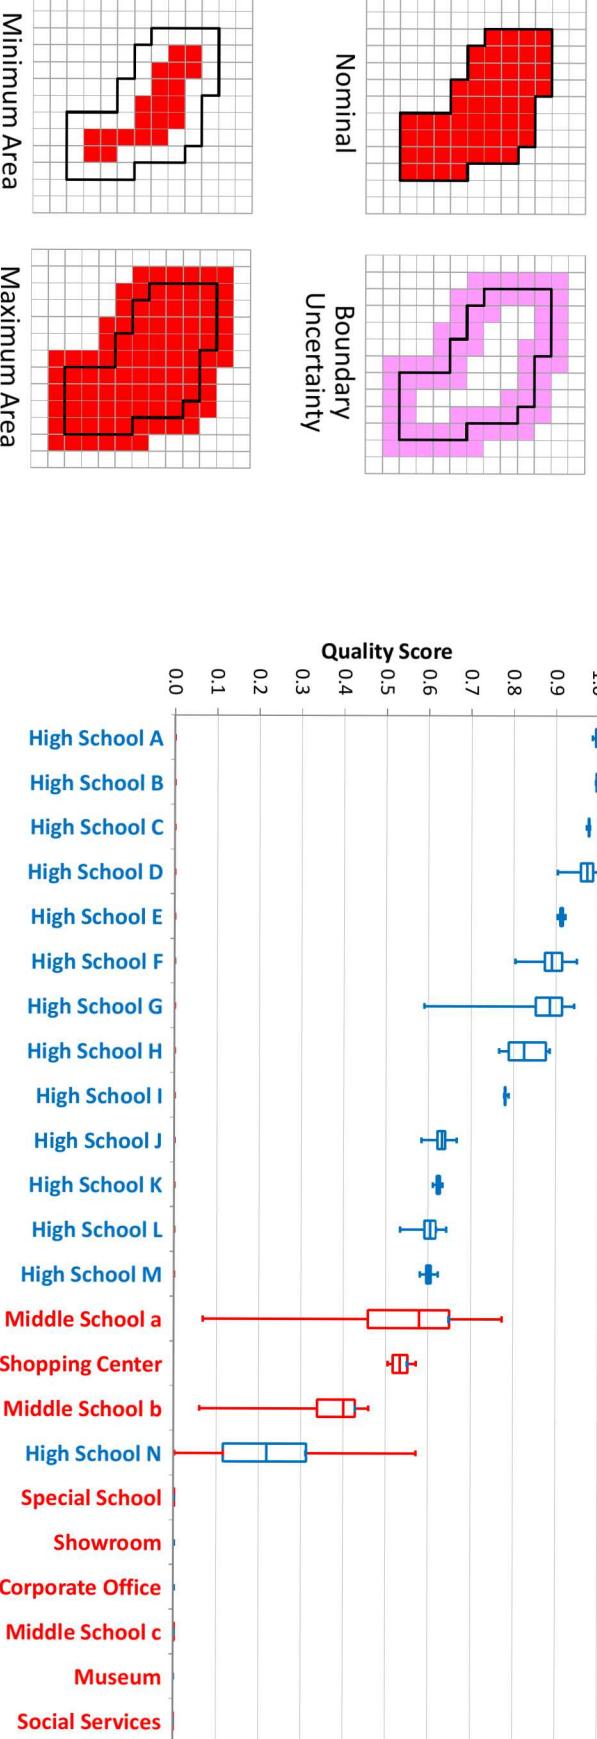
Example output:



Shape error model:



Results:



Sandia PANTHER Project

Goal:

Develop tools to aid understanding of voluminous remote sensor data, focusing on activity.

Main pillars:

- SAR image pre-processing (Sensor Exploitation).
- Signature search and trajectory analysis (Discrete Analytics).
- Human factors studies (Human Analytics).

Goal is to increase human analyst understanding and situational awareness.

Notes:

- Kristina Czuchlewski is the Principle Investigator.
- Supported by the Sandia LDRD office as a Grand Challenge – a substantial internal investment.
- PANTHER's signature search approach employs geospatial-temporal semantic graphs to analyze activity patterns.
- The projects are complementary: The NA-22 project focuses on durable objects, while PANTHER focuses on activity.

Why Activity Analysis?

Activity can help understand what's happening.

Successful?



Successful?



On Sunday:

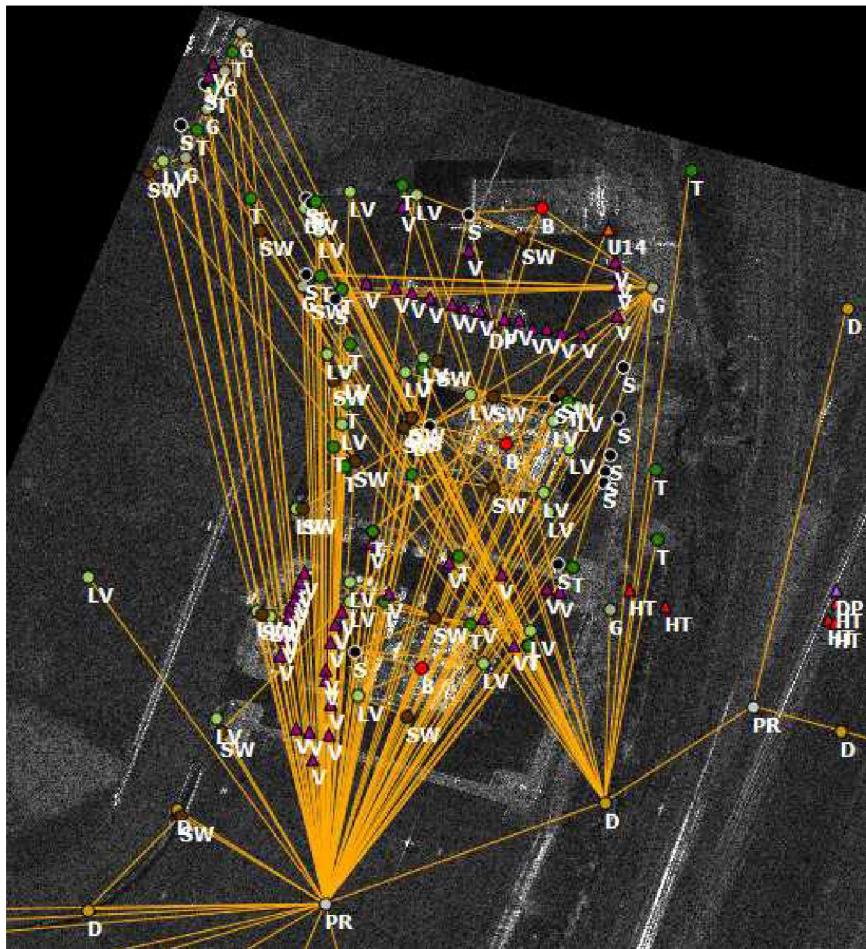


On Sunday:



Example: Activity Search

Geospatial-temporal semantic graph:



Match results:



Goal is to find large buildings with substantial nearby vehicle activity, indicating a possible active business.

Note: Based on hand-annotated primitive features.

Overview

- Motivation.
- Computation.
- Example Results.
- Extensions.
- ■ Discussion.

Overview

- Motivation.
- Computation.
- Example Results.
- Extensions.
- Discussion.

→ ▪ Recap.

- Application.
- Evaluation.
- Data supply.

Recap

We have demonstrated:

- Imagery + LiDAR + GIS \Rightarrow Land cover [O’Neil-Dunne 2012].
- Sequence \Rightarrow Geospatial-temporal graph.
- Spatial search: power plants, refineries, high schools...
- Spatial-temporal search: change, construction complexes...
- Over a wide area (2,067 km², 135 billion pixels, 3.6 million graph nodes).
- Multi-modality.

We have NOT shown:

- Continent-scale robust image supply and pre-processing.
- Recognition scope.
- Complex temporal analysis.
- Open issues:
Multiple hypotheses, match ambiguity, scale.

Overview

- Motivation.
- Computation.
- Example Results.
- Extensions.
- Discussion.
 - Recap.
- ▪ Application.
- Evaluation.
- Data supply.

Searching for a Site of Interest

- We have seen how this approach can solve a variety of search problems.
- Will it find proliferation sites of interest?
- We conducted a first experiment.

Steps:

1. Research sites, select candidates (LANL).
2. Obtain imagery and LiDAR (LLNL).
3. Image processing to find land cover (LLNL).
4. Construct geospatial-temporal graph (SNL).
5. Construct ontology, generate query (LANL).
6. Use query to search graph (SNL).

Results:

1. Code found site. ✓'s
2. Code did not find site in wide area where it was absent. ✓'s

This problem motivates an interest in global imagery.

Overview

- Motivation.
- Computation.
- Example Results.
- Extensions.
- Discussion.
 - Recap.
 - Application.
 - ▪ Evaluation.
 - Data supply.

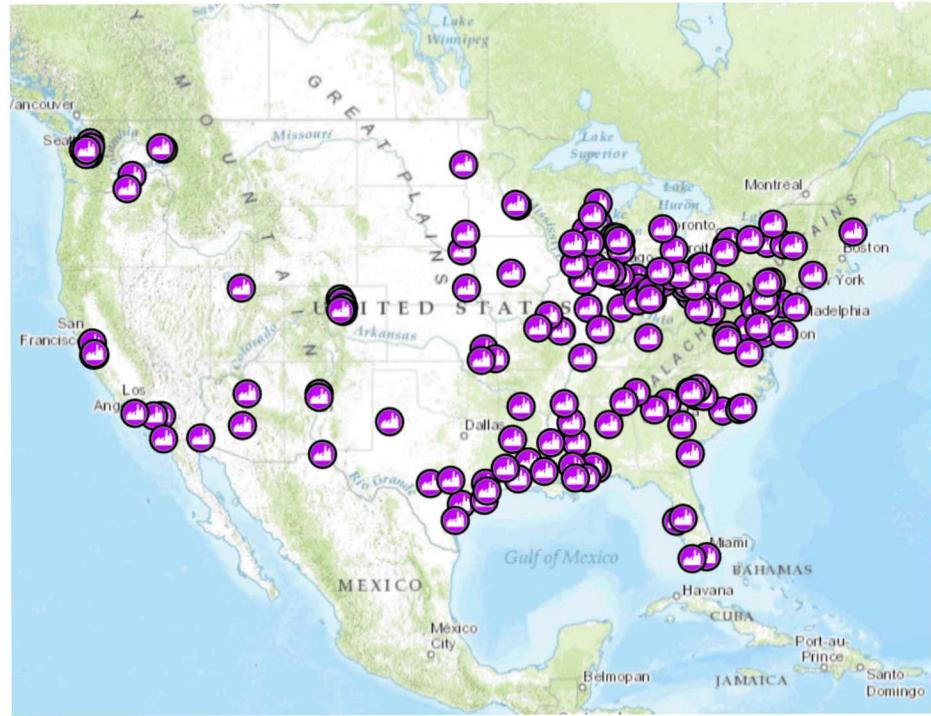
Evaluation

We are developing our assessment strategy:

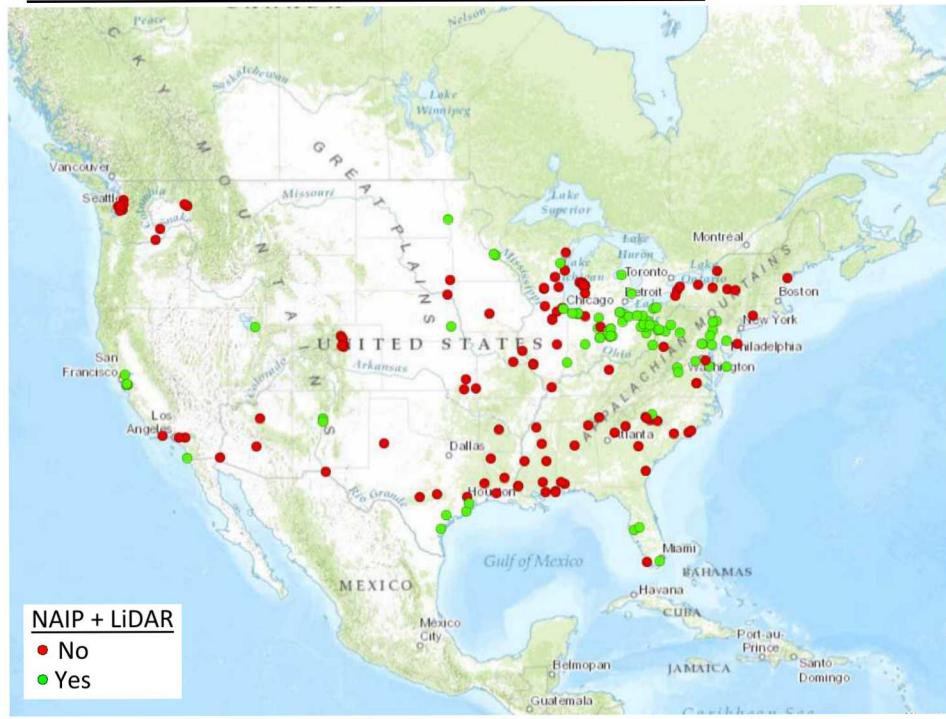
- By *assessment*, we mean essentially verification and validation (V&V).
 - Verification – “Are we building the thing right?”
 - Validation – “Are we building the right thing?”
- In research and development (R&D), requirements and metrics for evaluation evolve as the needs of stakeholders become clearer.
- Thus our strategy:
 1. Define algorithm requirements.
 2. Define metrics that quantify how well requirements are met.
 3. Evaluate performance wrt requirements.
 4. Determine if requirements or metrics require revision.
 5. Repeat until further changes are insignificant.
- Two phases:
 - A. Components: Image processing, ontology, graph search engine.
 - B. End-to-end: Probability of detection, false alarm rate, confidence levels.

UVM Analysis of Benchmark Sites

All sites:



80 locations with both NAIP and LiDAR:



We are aiming for a single rule set
that will solve all these sites.

Overview

- Motivation.
- Computation.
- Example Results.
- Extensions.
- Discussion.
 - Recap.
 - Application.
 - Evaluation.
- ▪ Data supply.

A Possible Data Supply Pipeline

1. Wide-area collect.
2. Orthorectification.
3. Atmospheric correction.
4. 3d stereo.
5. Land cover.
6. (Count cars.)
7. Store, setup provenance metadata.
8. Graph construction, change analysis.
9. Micro-search, with writeback.
10. RSS feed of changes of interest.
11. Ready for user interface.

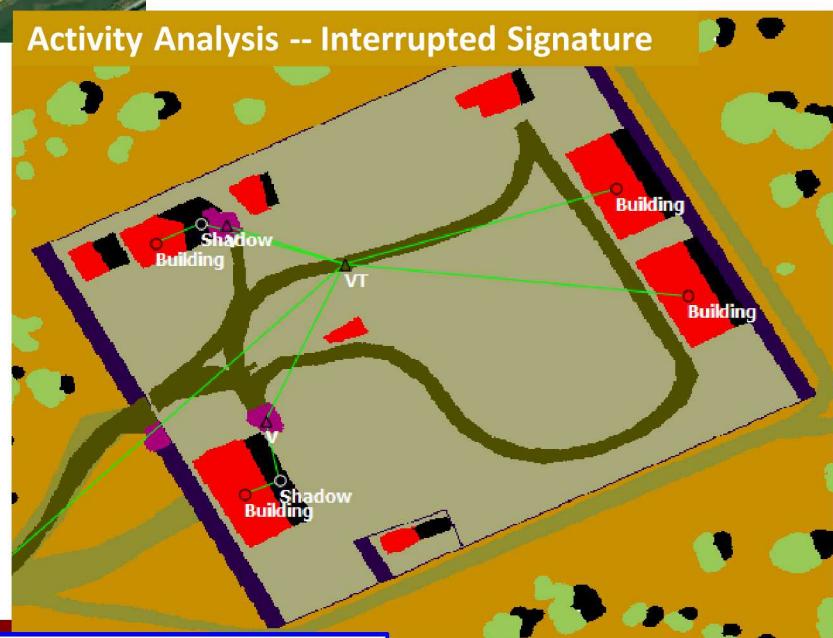
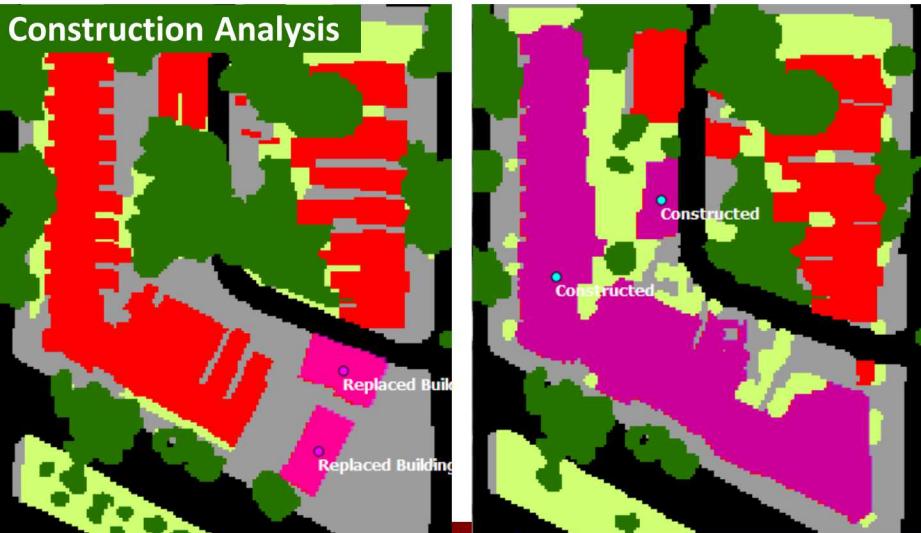
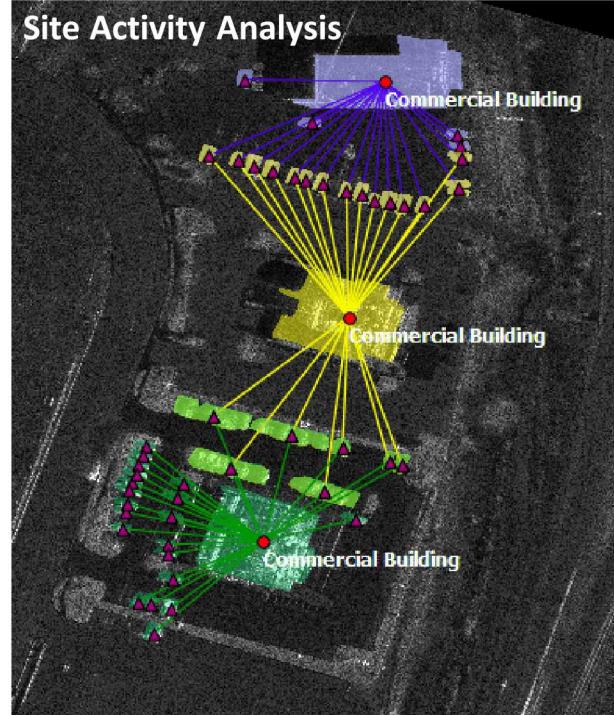
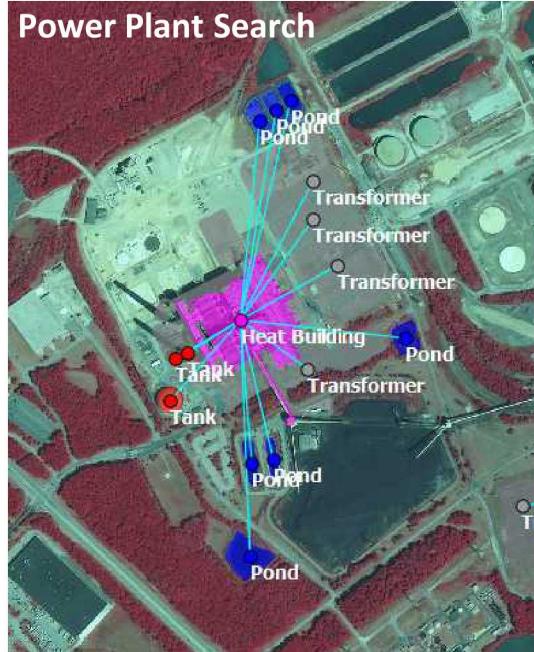
Questions:

- Feasible? Automatic? Reliable? World-wide?
- Are you thinking something different?

We're very interested
in your thoughts!

BACKUP SLIDES

Diversity of Problems



All of these were solved by the same code.

Search Run Times

As of January 2014:

	<u>Tri-Region*</u>
Power Plant	15.0 hours
Refinery	8.9 hours

	<u>Washington, DC**</u>
New Complex	1.7 hours

Reported in
BigSpatial paper
(Nov 2014).

May 2014:

High School	<u>Anne Arundel***</u>
	11.6 hours

Improvement:

March 2015:

High School (batch)	<u>Anne Arundel***</u>
High School (interactive)	0.1 hour
	0.03 hour (less than 2 minutes)

All times are single-thread,
single processor.

These algorithms are
well-suited to parallel
computation.

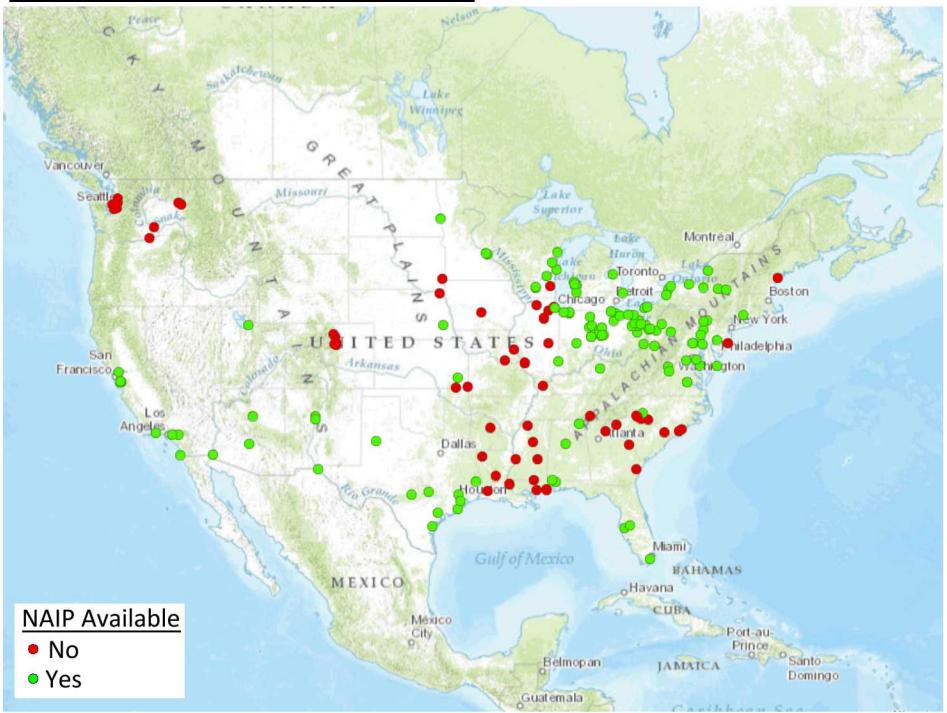
* Tri-Region: 2,067 km² area, over 9 billion land cover pixels, 3.5 million features.

** Washington, DC: 177 km² area, almost 1 billion land cover pixels, 1.3 million features, two times (2006, 2011).

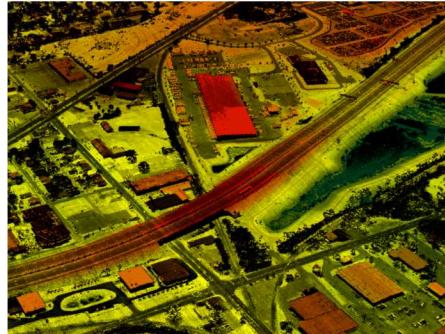
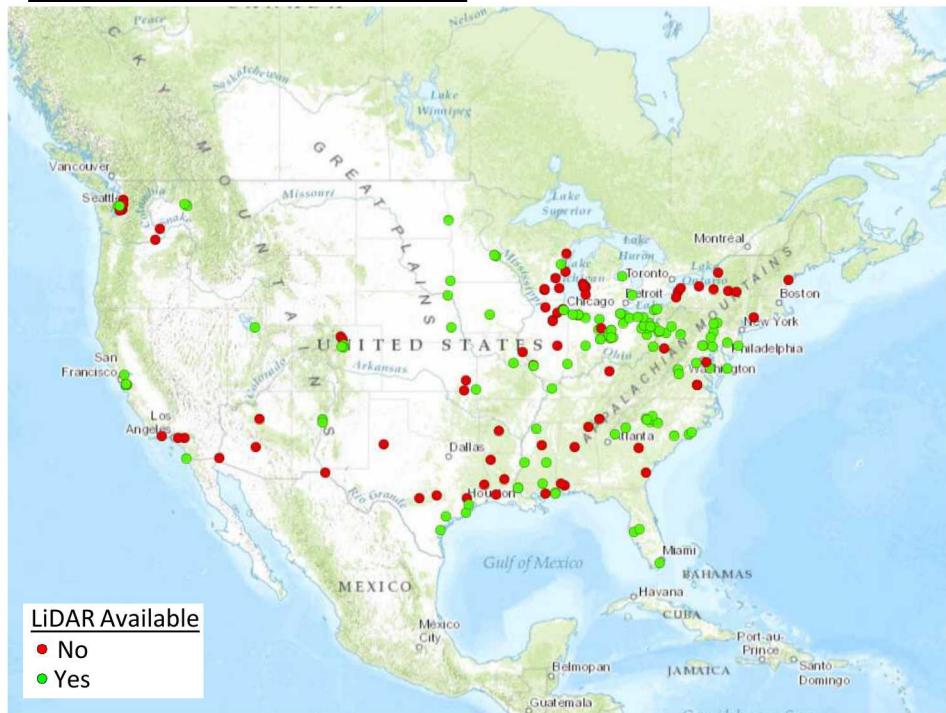
*** Anne Arundel County: 1,523 km² area, over 4 billion land cover pixels, 1.2 million features.

UVM: Survey of Supporting Data

124 locations with RGB+IR:



109 locations with LiDAR:



Simple Change Example

Resulting land cover:



2006



2011

Simple Change Example

(NA-22 Project)

Input data:



2006 Quickbird



2008 LiDAR



2011 NAIP

Image/Landcover Discrepancy

Input data:



The 2006 land cover shows two previous buildings, but the 2006 image shows these already removed.

Reason:

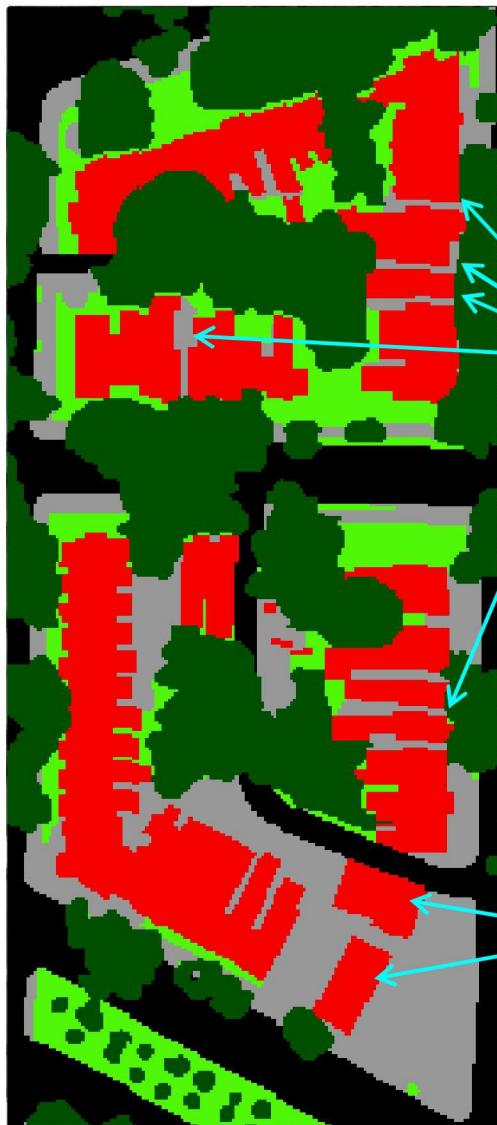
The land cover corresponds to a slightly different time than the 2006 Quickbird image.

Resolution:

We will treat the land cover data as the authority on the state of the world in 2006.

Geometric Change Analysis

Input land cover:



split

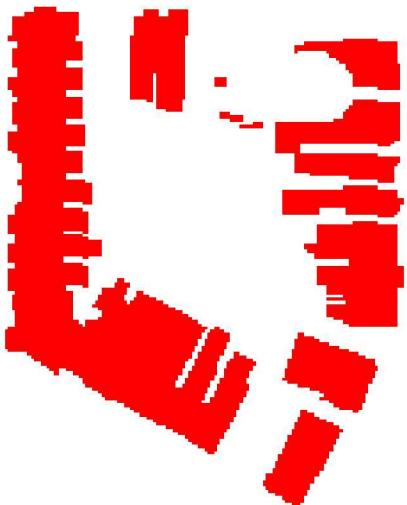
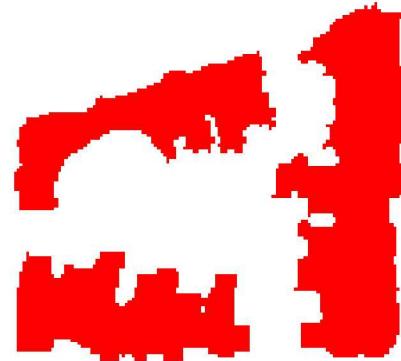
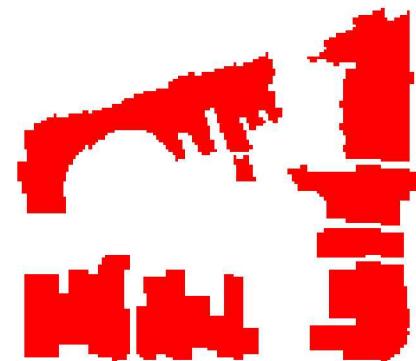
removed

new

added

Geometric Change Analysis

Geometric subtraction:

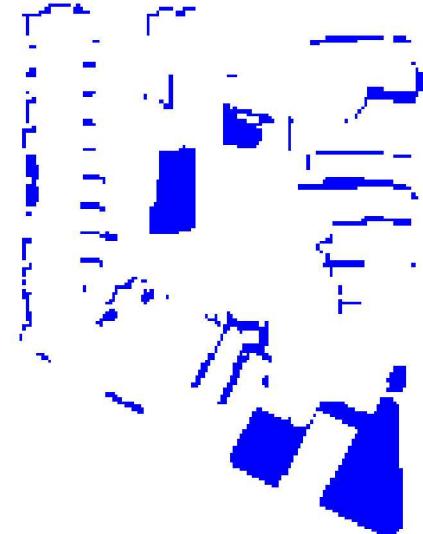
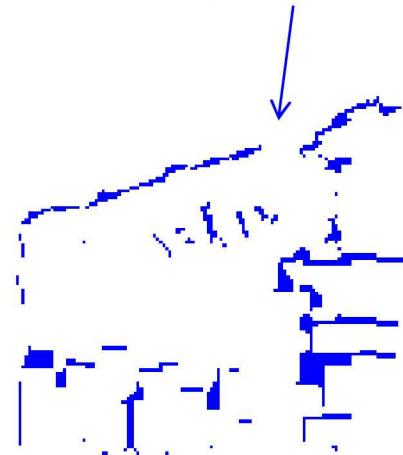


2006 Buildings

2011 Buildings

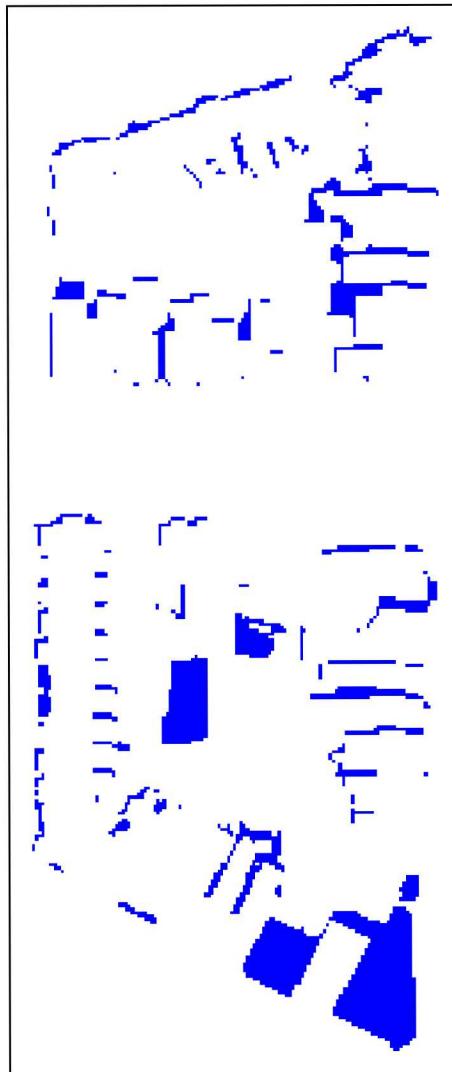
2011 minus 2006

Ghost boundaries result from sensor noise and registration error.

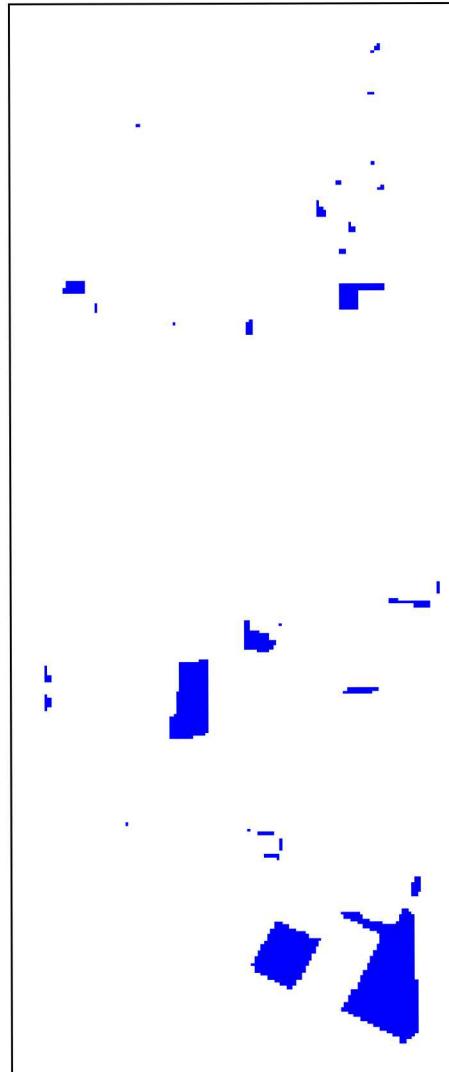


Geometric Change Analysis

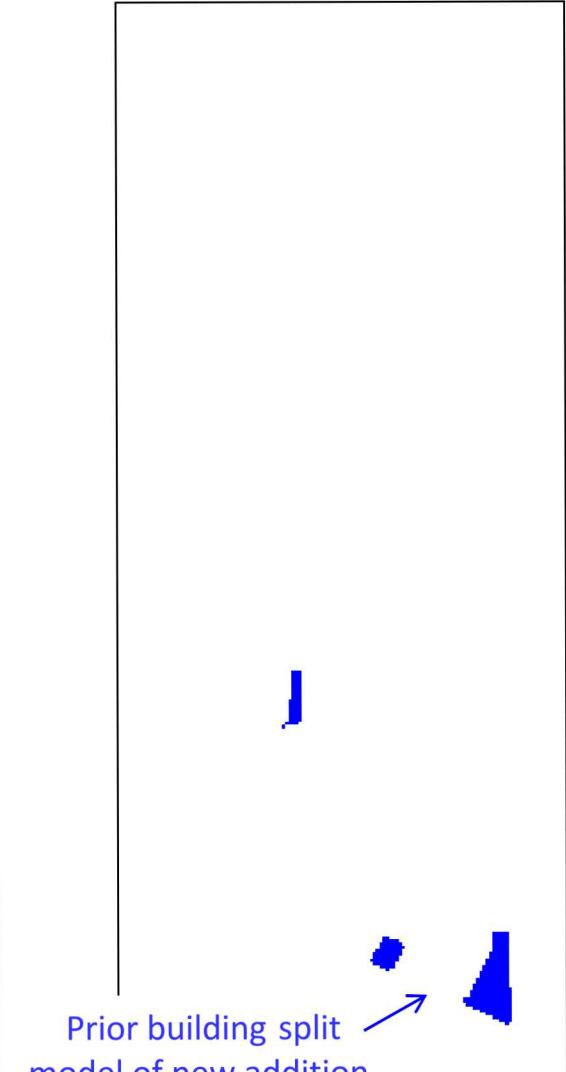
Shrinking to remove ghost boundaries:



2011 minus 2006



Shrink 1 Pixel

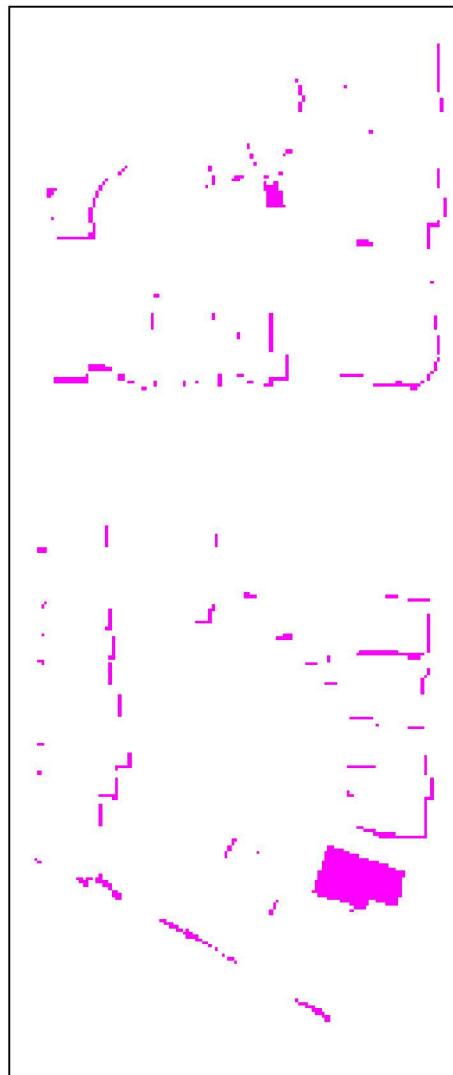


Prior building split
model of new addition.

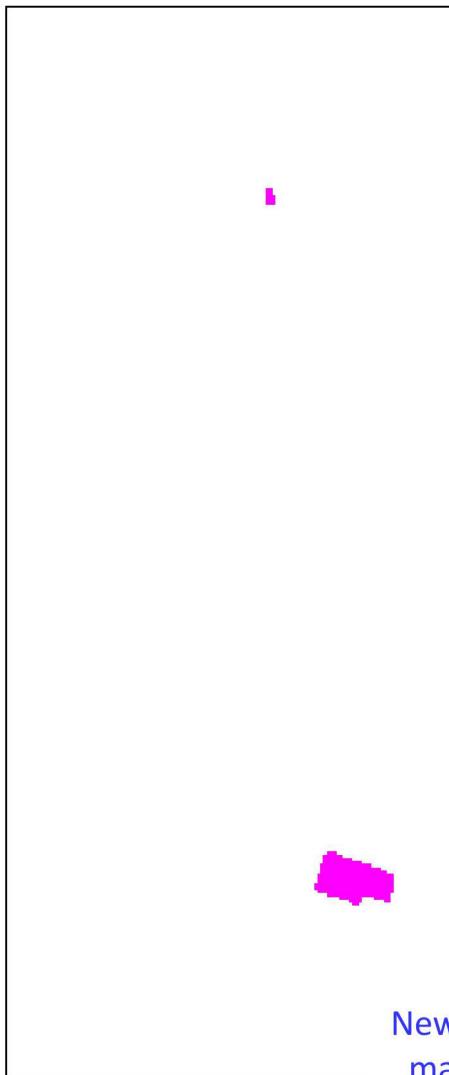
Shrink 4 Pixels

Geometric Change Analysis

Finding lost buildings (old minus new):



2006 minus 2011



Shrink 1 Pixel



Shrink 4 Pixels

New building extension
masked lost building.

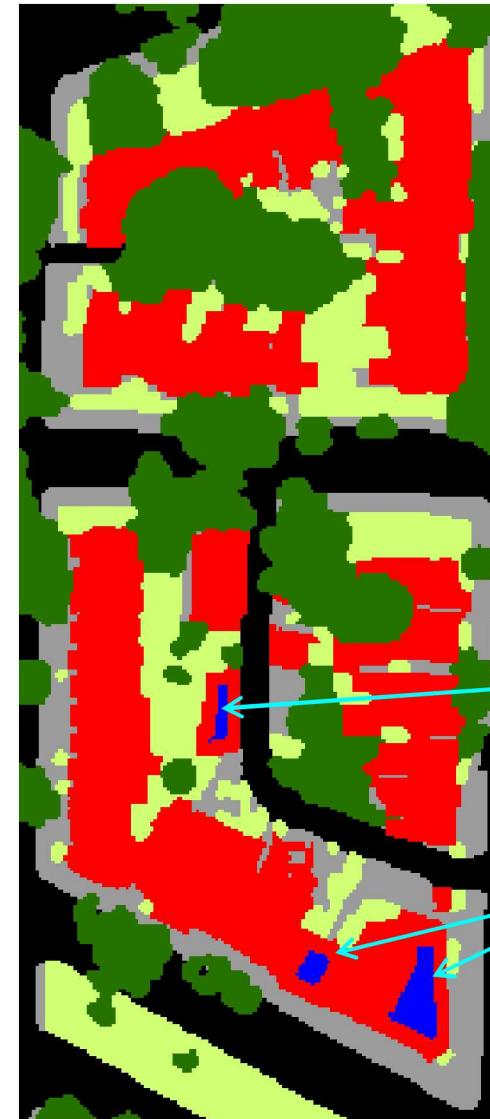


Geometric Change Analysis Results

Input land cover:



2006



2011

Change results
do not match
full shape.

removed
missed
this one

new
added
split in two;
misses interior
change section.

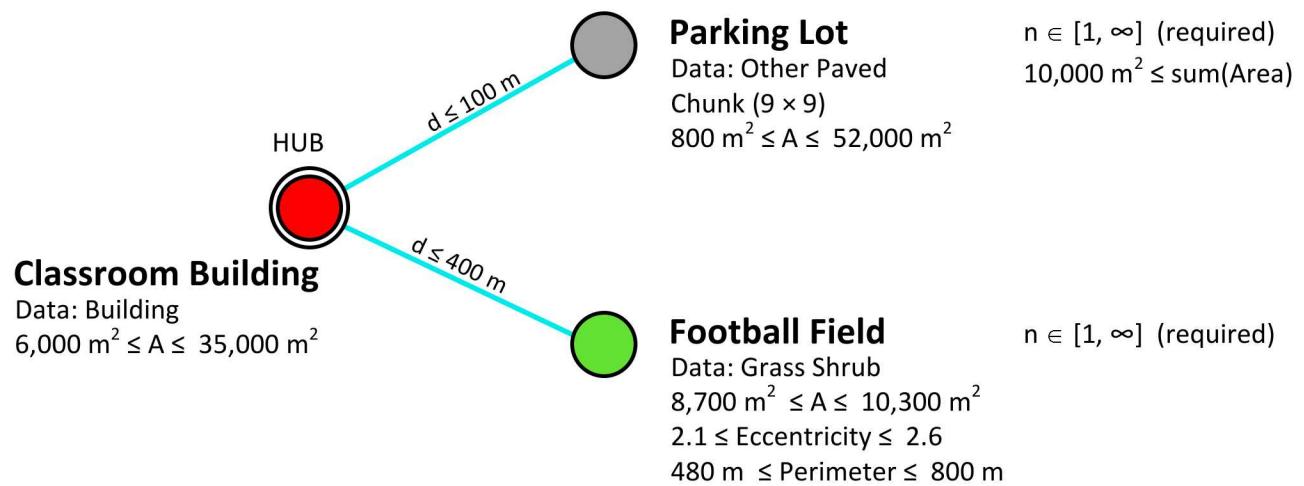
Simple High School Search

Query specification:

A high school is a classroom building with a parking lot and football field.



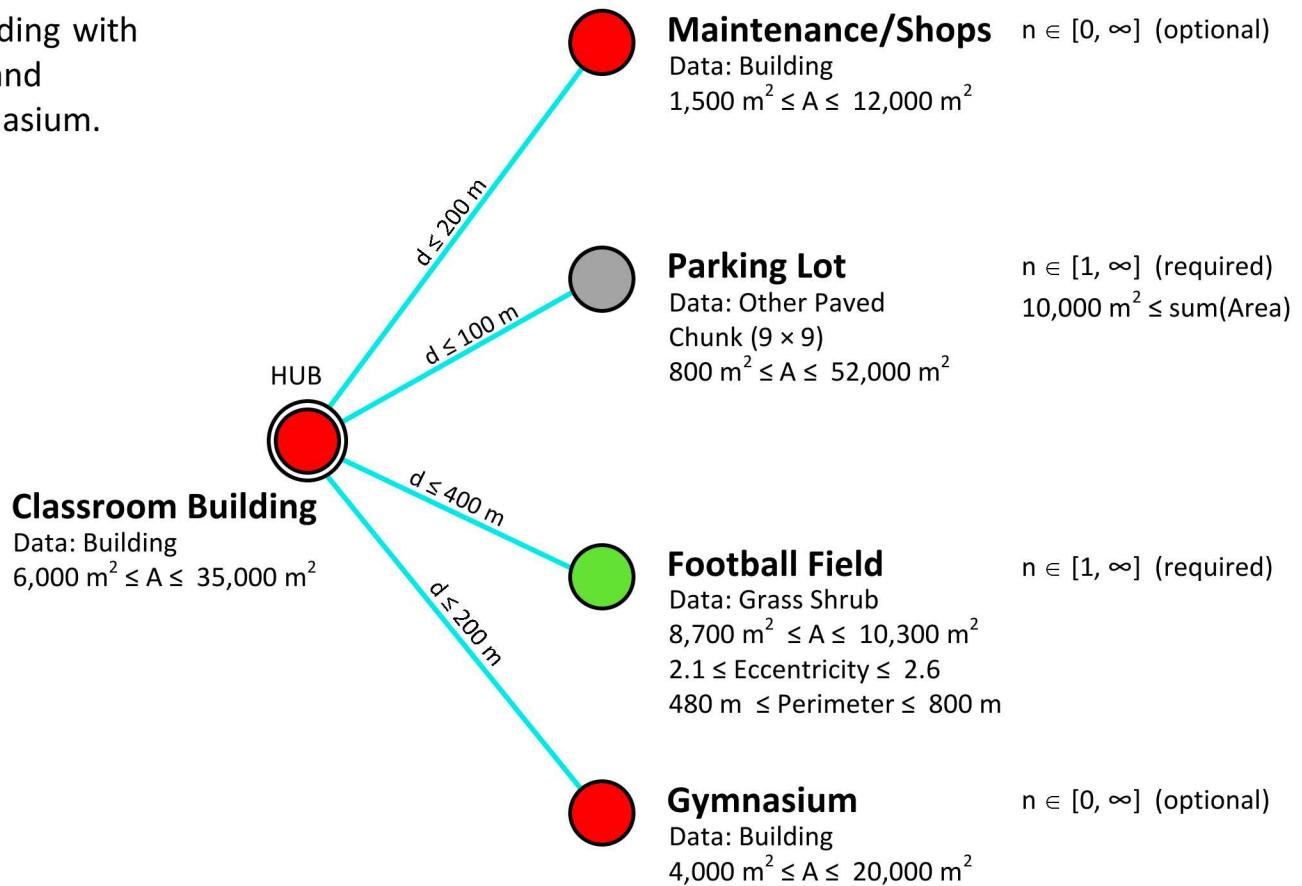
Northeast High School



Simple High School, Extended

Query specification:

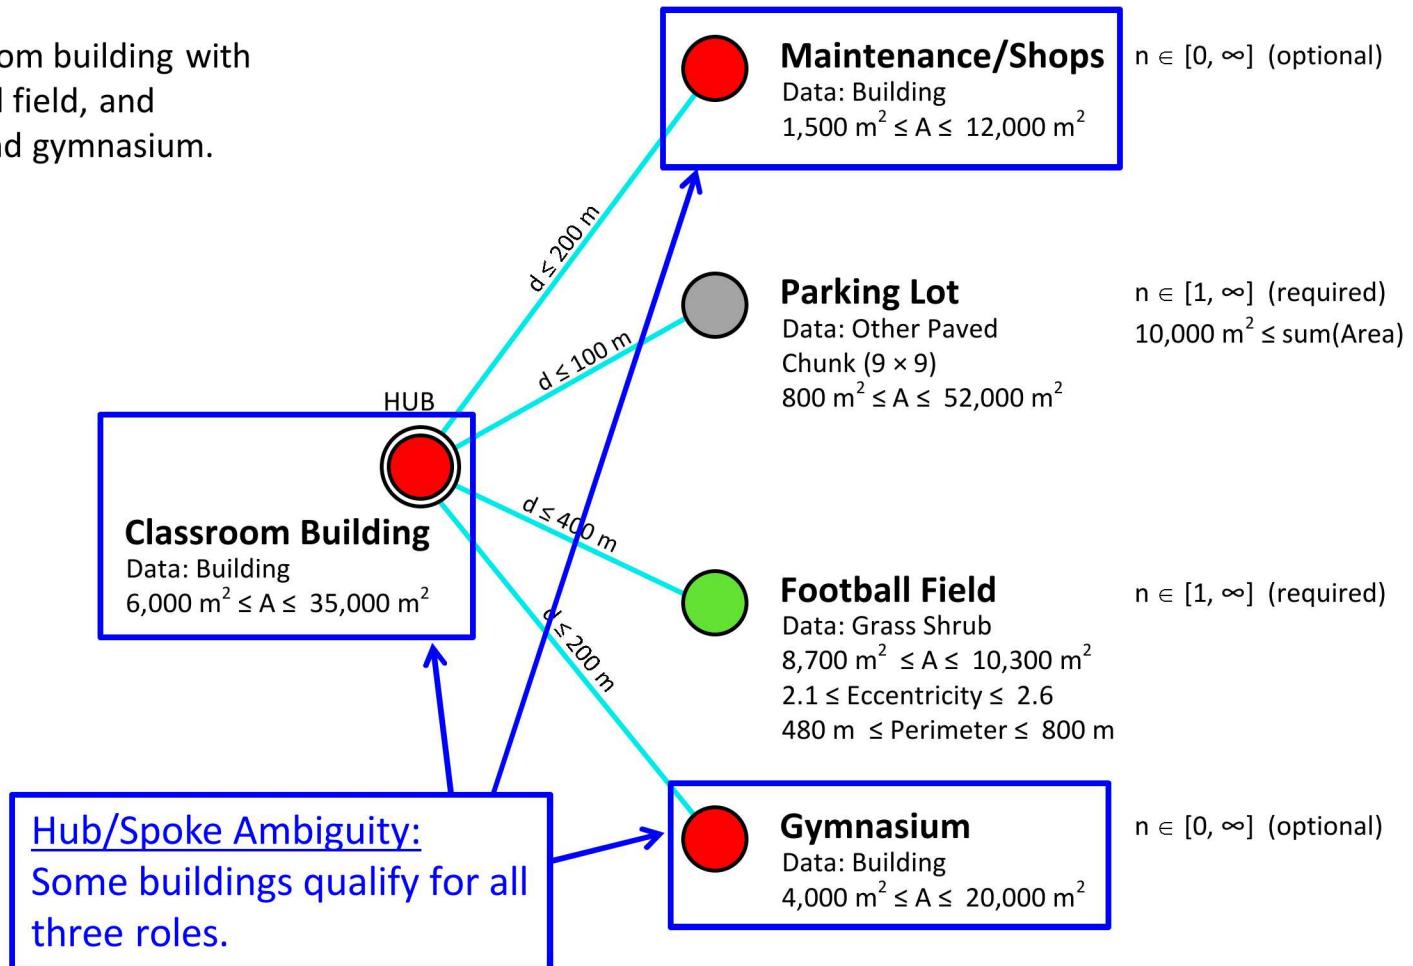
A high school is a classroom building with a parking lot and football field, and optional maintenance and gymnasium.



Hub/Spoke Ambiguity

Query specification:

A high school is a classroom building with a parking lot and football field, and optional maintenance and gymnasium.

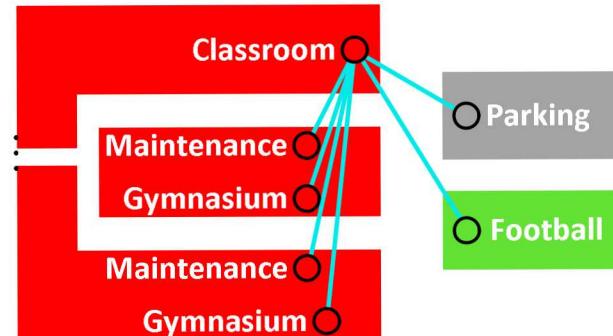


Ambiguous Hub/Spoke Search Result

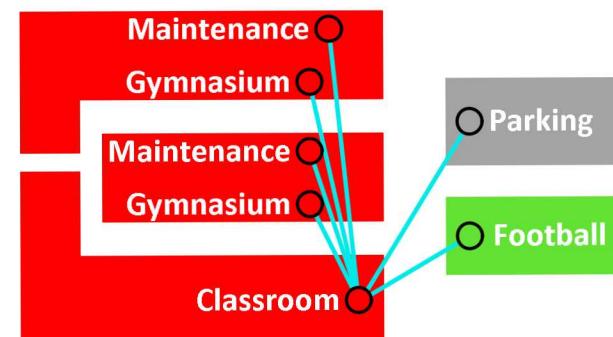
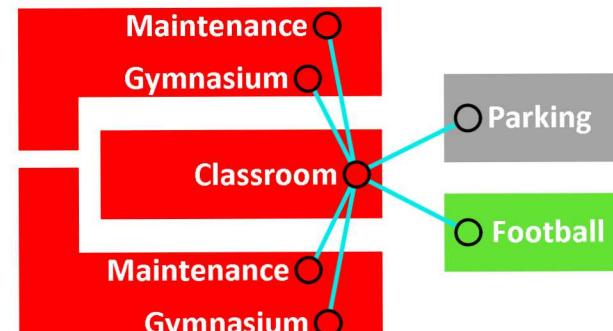
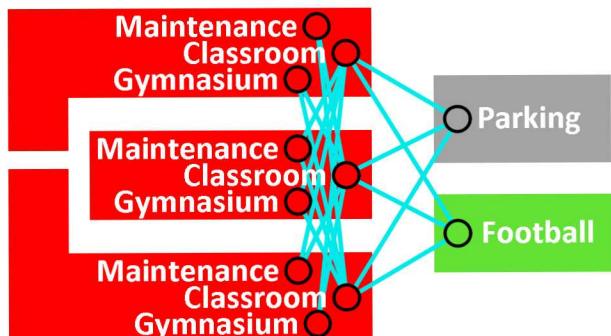
Scene:



Matches:



Search Graph:



PANTHER UI Rendering Query Results

View query results geo-spatially

Toggle map layers, including the graph nodes and edges, regions, and underlying source imagery for each result

Java World Wind

Layers

- MS Virtual Earth Aerial
- Bing Imagery
- Political Boundaries
- Place Names
- World Map
- Compass
- AnneArundelCounty 2007 Landcove

Fatal Accident_3

Java World Wind

Layers

- MS Virtual Earth Aerial
- Bing Imagery
- Political Boundaries
- Place Names
- World Map
- Compass
- AnneArundelCounty 2007 Landcove

Fatal Accident_2

Java World Wind

Layers

- MS Virtual Earth Aerial
- Bing Imagery
- Political Boundaries
- Place Names
- World Map
- Compass
- AnneArundelCounty 2007 Landcove

Selected Match Layers

- High Schools and Accidents_5
- Graph
- Regions
- Classroom Building_1
- 11/1/04 6:00 AM - 9/2/07 8:00 AM
 - Image - AnneArundel_HS_North
 - DSM - AnneArundel_HS_North
 - DEM - AnneArundel_HS_North
 - nDSM - AnneArundel_HS_North
- Parking Lot_1
- Football Field_1
- Fatal Accident_1
- Fatal Accident_2
- Fatal Accident_3
- Fatal Accident_4
- Fatal Accident_5

Parking Lot_1

Classroom Building_1

Football Field_1

Fatal Accident_1

Fatal Accident_2

Fatal Accident_3

Fatal Accident_4

Fatal Accident_5

Off Globe

Altitude 1 km

Off Globe

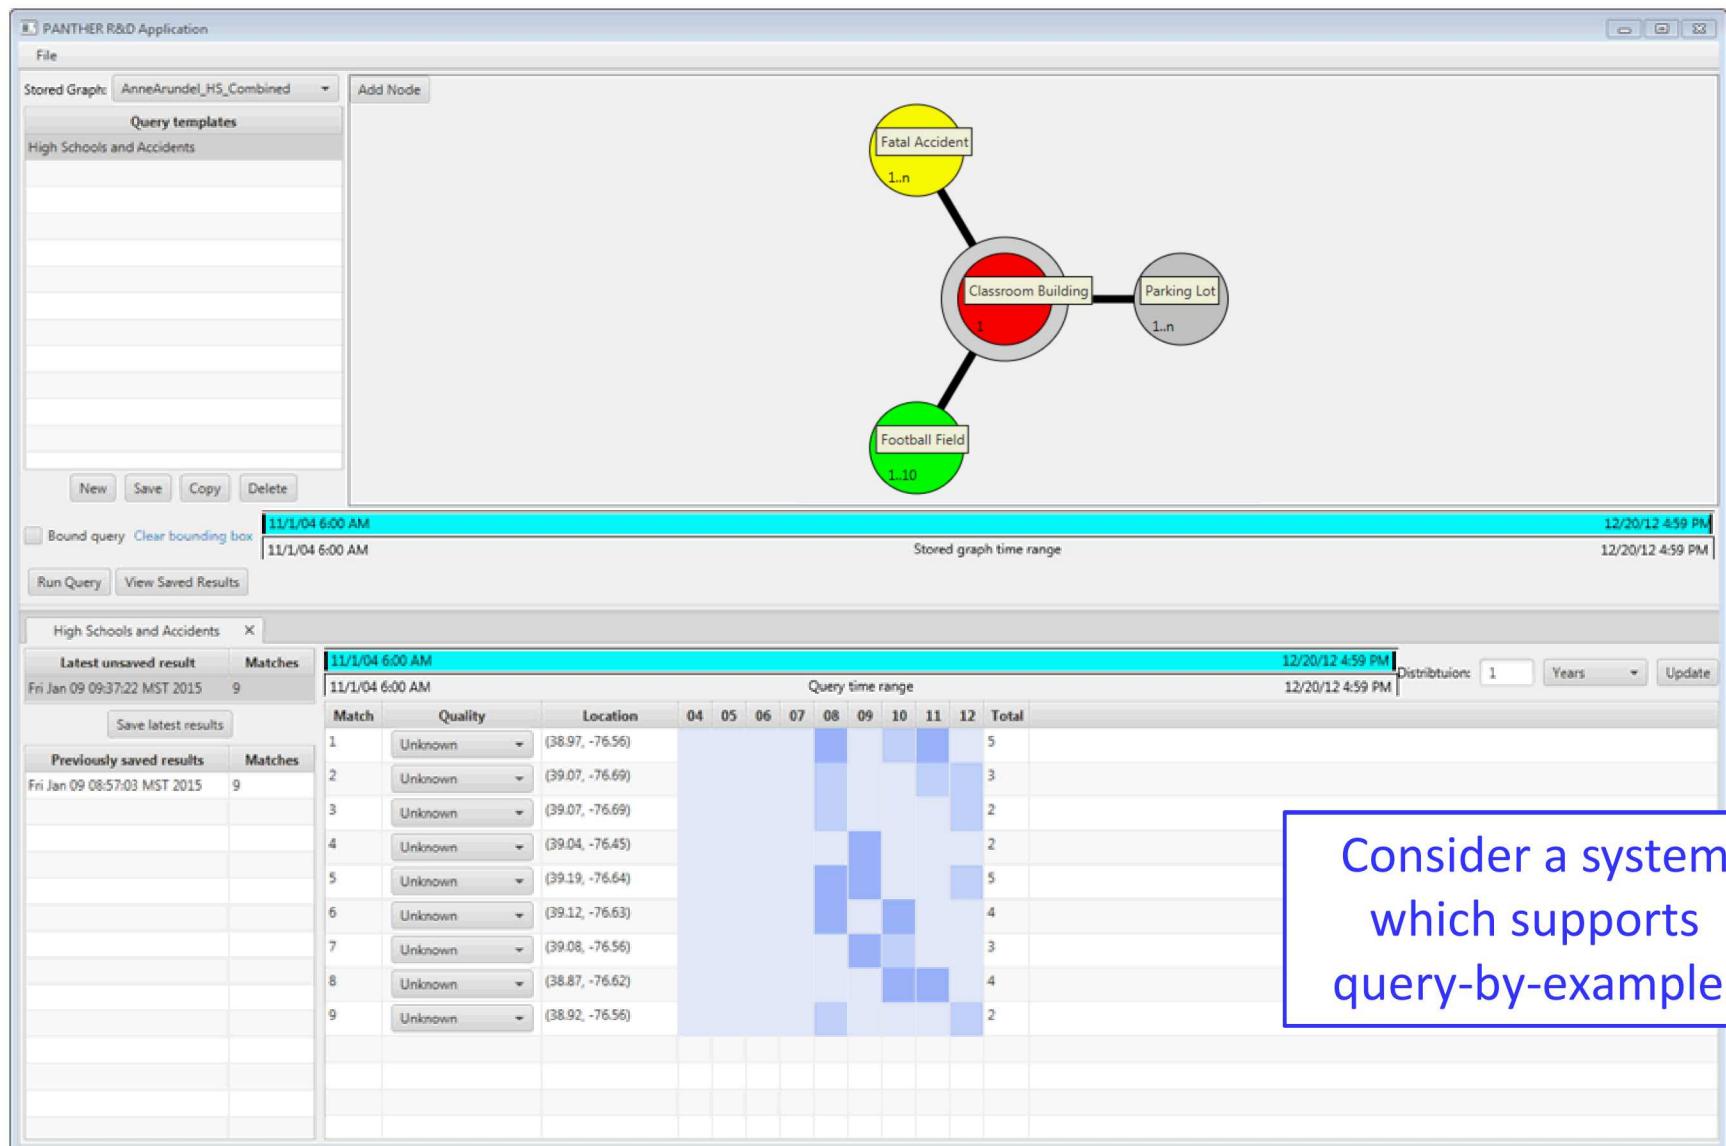
Altitude 2 km

Off Globe

200 m

PANTHER user interface by Jamie Coram, David Perkins, and Dan Morrow (SNL).

PANTHER User Interface



PANTHER user interface by Jamie Coram, David Perkins, and Dan Morrow (SNL).

Match Quality Scoring

Motivation:

- Some matches are more relevant than others.
- False negatives should be avoided.
- Widening search bounds captures false negatives, but also increases false positives.
- Quality scoring moves best matches to front, poor false positives to back of list.
(Similar to web search*)

* Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems 30 (1998) 107- 117, 30:107–117, 1998.*